arXiv:1512.06246v1 [cs.DB] 19 Dec 2015

Parallel-Correctness and Containment for Conjunctive Queries with Union and Negation

Gaetano Geck
TU Dortmund University

Bas Ketsman
Hasselt University

Frank Neven
Hasselt University

Thomas Schwentick
TU Dortmund University

Abstract

Single-round multiway join algorithms first reshuffle data over many servers and then evaluate the query at hand in a parallel and communication-free way. A key question is whether a given distribution policy for the reshuffle is adequate for computing a given query, also referred to as parallel-correctness. This paper extends the study of the complexity of parallel-correctness and its constituents, parallel-soundness and parallel-completeness, to unions of conjunctive queries with and without negation. As a by-product it is shown that the containment problem for conjunctive queries with negation is coNEXPTIME-complete.

1 Introduction

Motivated by recent in-memory systems like Spark [1] and Shark [20], Koutris and Suciu introduced the massively parallel communication model (MPC) [14] where computation proceeds in a sequence of parallel steps each followed by global synchronisation of all servers. Of particular interest in the MPC model are queries that can be evaluated in one round of communication [9]. In its most naïve setting, a query $\mathcal{Q}$ is evaluated by reshuffling the data over many servers, according to some distribution policy, and then computing $\mathcal{Q}$ at each server in a parallel but communication-free manner. A notable family of distribution policies is formed within the Hypercube algorithm [2, 9, 11]. A property of Hypercube distributions is that for any instance I, the central execution of $\mathcal{Q}(I)$ always equals the union of the evaluations of $\mathcal{Q}$ at every computing node (or server). The latter guarantees the correctness of the distributed evaluation for any conjunctive query by the Hypercube algorithm.

Ameloot et al. [5] introduced a general framework for reasoning about one-round evaluation algorithms under *arbitrary* distribution policies. They introduced *parallel-correctness* as a property of a query w.r.t. a distribution policy which states that central execution always equals distributed execution, that is, equals the union of the evaluations of the query at each server under the

given distribution policy. One of the main results of [5] is that deciding parallel-correctness for conjunctive queries (CQs) is Π_2^P-complete under arbitrary distribution policies. The upper bound follows rather directly from a semantical characterisation of parallel-correctness in terms of properties of minimal valuations. Specifically, it was shown that a conjunctive query is parallel-correct w.r.t. a distribution policy, if the distribution policy sends for every minimal valuation its required facts to at least one node.

As union and negation are fundamental operators, we extend in this paper the study of parallel-correctness to unions of conjunctive queries (UCQ), conjunctive queries with negation ($\text{CQ}^\neg$) and unions of conjunctive queries with negation ($\text{UCQ}^\neg$). In fact, we study two addional but related notions: parallel-soundness and parallel-completeness. While parallel-correctness implies equivalence between centralised and distributed execution, parallel-soundness (respectively, parallel-completeness) requires that distributed execution is contained in (respectively, contains) centralised execution. Of course, parallel-soundness and parallel-completeness together are equivalent to parallel-correctness. Furthermore, since all monotone queries are parallel-sound, on this class parallel-correctness is equivalent to parallel-completeness.

We start by investigating parallel-correctness for UCQ. Interestingly, for a UCQ to be parallel-correct under a certain distribution policy it is not required that every disjunct is parallel-correct. We extend the characterisation for parallel-correctness in terms of minimal valuations for CQs to UCQs and thereby obtain membership in Π_2^P. The matching lower bound follows, of course, from the lower bound for CQs [5].

Next, we study parallel-correctness for (unions of) conjunctive queries with negation. Sadly, when negation comes into play, parallel-correctness can no longer be characterised in terms of properties of valuations. Instead our algorithms are based on counter-examples of exponential size, yielding coNEXPTIME upper bounds. It turns out that this is optimal, though, as our corresponding lower bounds show. The proof of the lower bounds comes along an unexpected route: we exhibit a reduction from query containment for $\text{CQ}^\neg$ to parallel-correctness of $\text{CQ}^\neg$ (and its two variants) and show that query containment for $\text{CQ}^\neg$ is coNEXPTIME-complete. This is considerably different from what we thought was folklore knowledge of the community. Indeed, the Π_2^p-completeness result for query containment for $\text{CQ}^\neg$ mentioned in [18] only seems to hold for fixed database schemas (or a fixed arity bound, for that matter). We note that Mugnier et al. [16] provide a Π_2^p upper bound proof for $\text{CQ}^\neg$ containment and explicitly mention that it holds under the assumption that the arity of predicates is bounded by a constant. Altogether, parallel-correctness (and its variants) for (unions of) conjunctive queries with negation is thus complete for coNEXPTIME.

Finally, a natural question is how the high complexity of parallel-correctness in the presence of negation can be lowered. We identify two cases in which the complexity drops. More specifically, the complexity decreases from coNEXPTIME to Π_2^p if the database schema is fixed or the arity of relations is bounded, and to coNP for unions of *full* conjunctive queries with negation. In the latter case,

we again employ a reduction from containment of full conjunctive queries (with negation) and obtain novel results on the containment problem in this setting as well. All upper bounds hold for queries with inequalities.

Outline. This paper is further organised as follows. In Section 2, we discuss related work. In Section 3, we introduce the necessary definitions. We address parallel-correctness for unions of conjunctive queries in Section 4. We consider containment of conjunctive queries with negation in Section 5 and parallel-correctness together with its variants in Section 6. We discuss the restriction to full conjunctive queries in Section 7. We conclude in Section 8.

2 Related work

As mentioned in the introduction, Koutris and Suciu introduced the massively parallel communication model (MPC) [14]. A key property is that computation proceeds in a sequence of parallel steps, each followed by global synchronisation of all computing nodes. In this model, evaluation of conjunctive queries [8, 14] and skyline queries [4] has been considered. Beame, Koutris and Suciu [9] proved a matching upper and lower bound for the amount of communication needed to compute a full conjunctive query without self-joins in one communication round. The upper bound is provided by a randomised algorithm called *Hypercube* which uses a technique that can be traced back to Ganguly, Silberschatz, and Tsur [13] and is described in the context of map-reduce by Afrati and Ullman [2].

Ameloot et al. [5] introduced a general framework for reasoning about one-round evaluation algorithms under *arbitrary* distribution policies. They introduced the notion of *parallel-correctness* and proved its associated decision problem to be Π_2^p-complete for conjunctive queries. In addition, towards optimisation in MPC, they considered parallel-correctness *transfer*. Here, parallel-correctness transfers from $\mathcal{Q}$ to $\mathcal{Q}'$ when $\mathcal{Q}'$ is parallel-correct under every distribution policy for which $\mathcal{Q}$ is parallel-correct. The associated decision problem for conjunctive queries is shown to be Π_3^p-complete. In addition, some restricted cases (e.g., transferability under Hypercube distributions), are shown to be NP-complete.

Our definition of a distribution policy is borrowed from Ameloot et al. [6] (but already surfaces in the work of Zinn et al. [21]), where distribution policies are used to define the class of policy-aware transducer networks. The work by Ameloot et al. [7, 6] relates coordination-free computation with definability in variants of Datalog. One-round communication algorithms in MPC can be seen as very restrictive coordination-free computation.

The complexity of query containment for conjunctive queries is proved to be NP-complete by Chandra and Merlin [10]. Levy and Sagiv provide a test for query containment of conjunctive queries with negation [15] that involves exploring an exponential number of possible counter-example instances. In the context of information integration, Ullman [18] gives a comprehensive overview of query containment (with and without negation) and states the complexity of query containment for $\text{CQ}^\neg$ to be Π_2^p-complete. As mentioned in the introduc-

tion, the latter apparently only holds when the database schema is fixed or the arity of relations is considered to be bounded. A proof for the Π_2^p-lowerbound is given by Farré et al. [12]. Based on [15], Wei and Lausen [19] study a method for testing containment that exploits containment mappings for the positive parts of queries, and additionally provide a characterisation for UCQ$^\neg$ containment.

3 Definitions

3.1 Queries and instances

We assume an infinite set **dom** of data values that can be represented by strings over some fixed alphabet. By $\mathbf{dom}_n$ we denote the set of data values represented by strings of length at most n. A *database schema* $\mathcal{D}$ is a finite set of relation names R, each with some arity $ar(R)$. We also write $R^{(k)}$ as a shorthand to denote that R is a relation of arity k. We call $R(\mathbf{t})$ a *fact* when R is a relation name and $\mathbf{t}$ a tuple over **dom** of appropriate arity. We say that a fact $R(\mathbf{t})$ is *over* a database schema $\mathcal{D}$ if $R \in \mathcal{D}$. For a subset $U \subseteq \mathbf{dom}$ we write $facts(\mathcal{D}, U)$ for the set of possible facts over schema $\mathcal{D}$ and U and by $facts(\mathcal{D})$ we denote $facts(\mathcal{D}, \mathbf{dom})$. A *(database) instance* I over $\mathcal{D}$ is a finite set of facts over $\mathcal{D}$. By $adom(I)$ we denote the set of data values occurring in I. A *query* $\mathcal{Q}$ *over input schema* $\mathcal{D}_1$ *and output schema* $\mathcal{D}_2$ is a generic mapping from instances over $\mathcal{D}_1$ to instances over $\mathcal{D}_2$. Genericity means that for every permutation π of **dom** and every instance I, $\mathcal{Q}(\pi(I)) = \pi(\mathcal{Q}(I))$. We say that $\mathcal{Q}$ is *contained* in $\mathcal{Q}'$, denoted $\mathcal{Q} \subseteq \mathcal{Q}'$ iff for all instances I, $\mathcal{Q}(I) \subseteq \mathcal{Q}'(I)$.

3.2 Unions of conjunctive queries with negation

Let **var** be an infinite set of variables, disjoint from **dom**. An *atom* over schema $\mathcal{D}$ is of the form $R(\mathbf{x})$, where R is a relation name from $\mathcal{D}$ and $\mathbf{x} = (x_1, \ldots, x_k)$ is a tuple of variables in **var** with $k = ar(R)$. A *conjunctive query* $\mathcal{Q}$ *with negation and inequalities* over input schema $\mathcal{D}$ is an expression of the form

$$T(\mathbf{x}) \leftarrow R_1(\mathbf{y_1}), \ldots, R_m(\mathbf{y_m}), \neg S_1(\mathbf{z_1}), \ldots, \neg S_n(\mathbf{z_n}), \beta_1, \ldots, \beta_p$$

where all $R_i(\mathbf{y}_i)$ and $S_i(\mathbf{z}_j)$ are atoms over $\mathcal{D}$, every β_i is an inequality of the form $s \neq s'$ where s, s' are distinct variables occurring in some $\mathbf{y}_i$ or $\mathbf{z}_j$, and $T(\mathbf{x})$ is an atom for which $T \notin \mathcal{D}$. Additionally, for safety, we require that every variable in $\mathbf{x}$ occurs in some $\mathbf{y}_i$ and that every variable occurring in a negated atom has to occur in a positive atom as well (*safe* negation). We refer to the *head atom* $T(\mathbf{x})$ as $head_\mathcal{Q}$, to the set $\{R_1(\mathbf{y_1}), \ldots, R_m(\mathbf{y_m}), S_1(\mathbf{z_1}), \ldots, S_n(\mathbf{z_n})\}$ as $body_\mathcal{Q}$, and to the set $\{\beta_1, \ldots, \beta_p\}$ as $ineq_\mathcal{Q}$. Specifically, we refer to $\{R_1(\mathbf{y_1}), \ldots, R_m(\mathbf{y_m})\}$ as the positive atoms in $\mathcal{Q}$, denoted $pos_\mathcal{Q}$, and to $\{S_1(\mathbf{z_1}), \ldots, S_n(\mathbf{z_n})\}$ as the *negated* atoms of $\mathcal{Q}$, denoted $neg_\mathcal{Q}$. We denote by $vars(\mathcal{Q})$ the set of all variables occurring in $\mathcal{Q}$. We refer to the class of conjunctive queries with negation and inequalities by $\mathbf{CQ}^{\neg,\neq}$, its restriction to queries without inequalities, without negated atoms, and without both by $\mathbf{CQ}^\neg$,

$\mathbf{CQ}^{\neq}$, and $\mathbf{CQ}$, respectively. As a shorthand we refer to queries from $\mathbf{CQ}^{\neg,\neq}$ as $\text{CQ}^{\neg,\neq}$s and similarly for the other classes.

A *pre-valuation* for a $\text{CQ}^{\neg,\neq}$ $\mathcal{Q}$ is a total function $V : \mathit{vars}(\mathcal{Q}) \to \mathbf{dom}$, which naturally extends to atoms and sets of atoms. It is *consistent* for $\mathcal{Q}$, if $V(\mathit{pos}_{\mathcal{Q}}) \cap V(\mathit{neg}_{\mathcal{Q}}) = \emptyset$, and $V(s) \neq V(s')$, for every inequality $s \neq s'$ of $\mathcal{Q}$, in which case it is called a valuation. Of course, for a conjunctive query without negated atoms and without inequalities, every pre-valuation is also a valuation. We refer to $V(\mathit{pos}_{\mathcal{Q}})$ as the facts *required* by V, and to $V(\mathit{neg}_{\mathcal{Q}})$ as the facts *prohibited* by V.

A valuation V *satisfies* $\mathcal{Q}$ on instance I if all facts required by V are in I while no fact prohibited by V is in I, that is, if $V(\mathit{pos}_{\mathcal{Q}}) \subseteq I$ and $V(\mathit{neg}_{\mathcal{Q}}) \cap I = \emptyset$. In that case, V *derives* the fact $V(\mathit{head}_{\mathcal{Q}})$. The *result of* $\mathcal{Q}$ *on instance* I, denoted $\mathcal{Q}(I)$, is defined as the set of facts that can be derived by satisfying valuations for $\mathcal{Q}$ on I.

A *union of conjunctive queries with negation and inequalities* is a finite union of $\text{CQ}^{\neg,\neq}$s. That is, $\mathcal{Q}$ is of the form $\bigcup_{i=1}^{n} \mathcal{Q}_i$ where all subqueries $\mathcal{Q}_1, \ldots, \mathcal{Q}_n$ have the same relation name in their head atoms. We assume disjoint variable sets among different disjuncts in $\mathcal{Q}$. That is, $\mathit{vars}(\mathcal{Q}_i) \cap \mathit{vars}(\mathcal{Q}_j) = \emptyset$ for $i \neq j$ and, in particular, $\mathit{vars}(\mathit{head}_{\mathcal{Q}_i}) \neq \mathit{vars}(\mathit{head}_{\mathcal{Q}_j})$. By $\mathit{varmax}(\mathcal{Q})$ we denote the maximum number of variables that occurs in any disjunct of $\mathcal{Q}$. By $\mathbf{UCQ}^{\neg,\neq}$ we denote the class of unions of conjunctive queries with negation and inequalities and its fragments are denoted correspondingly.

A $\text{CQ}^{\neg,\neq}$ is called *full* if all of its variables occur in its head. A $\text{UCQ}^{\neg,\neq}$ is *full* if all its subqueries are full.

The *result of* $\mathcal{Q}$ *on instance* I is $\mathcal{Q}(I) = \bigcup_{i=1}^{n} \mathcal{Q}_i(I)$. Accordingly, a mapping from variables to data values is a *valuation* for a $\text{UCQ}^{\neg,\neq}$ $\mathcal{Q}$ if it is a valuation for one of its subqueries.

3.3 Networks, data distribution, and policies

A *network* $\mathcal{N}$ is a nonempty finite set of values from $\mathbf{dom}$, which we call *(computing) nodes* (or servers). A *distribution policy* $\boldsymbol{P} = (U, \mathit{rfacts}_{\boldsymbol{P}})$ for a database schema $\mathcal{D}$ and a network $\mathcal{N}$ consists of a universe U and a total function $\mathit{rfacts}_{\boldsymbol{P}}$ that maps each node of $\mathcal{N}$ to a set of facts from $\mathit{facts}(\mathcal{D}, U)$. A node κ is *responsible for fact* $\boldsymbol{f}$ (under policy $\boldsymbol{P}$) if $\boldsymbol{f} \in \mathit{rfacts}_{\boldsymbol{P}}(\kappa)$. As a shorthand (and slight abuse of notation), we denote the set of nodes κ that are responsible for some given fact $\boldsymbol{f}$ by $\boldsymbol{P}(\boldsymbol{f})$. For a distribution policy $\boldsymbol{P}$ and an instance I over $\mathcal{D}$, let $\mathit{loc\text{-}inst}_{\boldsymbol{P},I}$ denote the function that maps each $\kappa \in \mathcal{N}$ to $I \cap \mathit{rfacts}_{\boldsymbol{P}}(\kappa)$, that is, the set of facts in I for which κ is responsible. We sometimes refer to a given instance I as the *global instance* and to $\mathit{loc\text{-}inst}_{\boldsymbol{P},I}(\kappa)$ as the *local instance at node* κ.

We note that for some facts from $\mathit{facts}(\mathcal{D}, U)$ there are no responsible nodes. This gives our framework some additional flexibility. However, it does not affect our results: in the lower bound proofs we only use distributions for which all facts from $\mathit{facts}(\mathcal{D}, U)$ have some responsible nodes. Each distribution policy

implicitly induces a network and each query implicitly defines a database (sub-) schema. Therefore, we often omit the explicit notation for networks and schemas.

Given some policy $\boldsymbol{P}$ that is defined over a network $\mathcal{N}$, the *result* $[\mathcal{Q}, \boldsymbol{P}](I)$ *of the distributed evaluation of a query* $\mathcal{Q}$ *on an instance* I *in one round* is defined as the union of the results of the query evaluated on each node's local instance. Formally,

$$[\mathcal{Q}, \boldsymbol{P}](I) \stackrel{\text{def}}{=} \bigcup_{\kappa \in \mathcal{N}} \mathcal{Q}\big(\mathit{loc\text{-}inst}_{\boldsymbol{P},I}(\kappa)\big).$$

In the decision problem for parallel correctness (to be formalised later), the input consists of a query $\mathcal{Q}$ and a distribution policy $\boldsymbol{P}$. However, it is not obvious how distribution policies should be specified. In principle, they could be defined in an arbitrary fashion, but it is reasonable to assume that given a potential fact $\boldsymbol{f}$, a node κ and a policy $\boldsymbol{P}$, it is not too hard to find out whether κ is responsible for $\boldsymbol{f}$ under $\boldsymbol{P}$.

For UCQ$^{\neq}$s, which are monotone, our complexity results are remarkably robust with respect to the choice of the representation of distribution policies. In fact, the complexity results coincide for the two extreme possible choices that we consider in this article. In the first case, distribution policies are specified by an explicit list of tuple-node-pairs, whereas in the second case the test whether a given node is responsible for a given tuple can be carried out by a non-deterministic polynomial-time algorithm. However, we do require that some bound n on the length of strings that represent node names and data values is given. Without such a restriction, no upper complexity bounds would be possible as nodes with names of super-polynomial length in the size of the input would not be accessible.

Considering queries with negated atoms, however, these two settings (seem to) differ, complexity-wise. Thus, we consider a third option, $\mathcal{P}_{\text{rule}}$, in which the universe U of a policy is explicitly enumerated and the responsibilities are defined by simple constraints (described below). The latter representation enjoys the same complexity properties as the full NP-test based case.

Now we give more precise definitions of classes of policies and their representations as inputs of algorithmic problems. As said before, policies $\boldsymbol{P} = (U, \mathit{rfacts}_{\boldsymbol{P}})$ from $\mathcal{P}_{\text{fin}}$ are specified by an explicit enumeration of U and of all pairs $(\kappa, \boldsymbol{f})$ where $\kappa \in \boldsymbol{P}(\boldsymbol{f})$. A policy $\boldsymbol{P} = (U, \mathit{rfacts}_{\boldsymbol{P}})$ from $\mathcal{P}_{\text{rule}}$ is given by an explicit enumeration of U and a list of *rules* of the form $\rho = (A, \kappa)$, where A is an atom with variables and/or constants from U, and a network node κ. The semantics of such a rule is as follows: for every substitution $\mu : \mathbf{var} \cup \mathbf{dom} \to \mathbf{dom}$ that maps variables to values from U and leaves constants from U unchanged, the node κ is responsible for the fact $\mu(A)$. A rule is a *fact rule* if its atom does not contain any variables, that is, $A = R(a_1, \ldots, a_n)$, where $a_1, \ldots, a_n \in U$. In particular, $\mathcal{P}_{\text{fin}} \subseteq \mathcal{P}_{\text{rule}}$.

Example 3.1. Let distribution policy $\boldsymbol{P}$ over schema $\{\texttt{Rel}^{(3)}\}$ and network $\{\kappa_1, \kappa_2\}$ be given by $U = \{1, \ldots, 10\}$ and the rules $\big(\texttt{Rel}(1, x, x), \kappa_1\big)$, $\big(\texttt{Rel}(2, x, y), \kappa_2\big)$.

On global instance $I = \{\texttt{Rel}(1,7,7), \texttt{Rel}(1,7,8), \texttt{Rel}(2,9,8), \texttt{Rel}(2,9,9)\}$, policy $\boldsymbol{P}$ induces local instances $\mathit{loc\text{-}inst}_{\boldsymbol{P},I}(\kappa_1) = \{\texttt{Rel}(1,7,7)\}$ and $\mathit{loc\text{-}inst}_{\boldsymbol{P},I}(\kappa_2) = \{\texttt{Rel}(2,9,8), \texttt{Rel}(2,9,9)\}$. □

The most general classes of policies allow to specify policies by means of a "test algorithm" with time bound ℓ^k, where ℓ is the length of the input and k some constant. Such an algorithm decides, for an input consisting of a node κ and fact $\boldsymbol{f}$, whether κ is responsible for $\boldsymbol{f}$.[1] A policy $\boldsymbol{P} = (U, \mathit{rfacts}_{\boldsymbol{P}})$ from $\mathcal{P}^k_{\text{npoly}}$ is specified by a pair $(n, \mathcal{A}_{\boldsymbol{P}})$, where n is a natural number in unary representation and $\mathcal{A}_{\boldsymbol{P}}$ is a non-deterministic algorithm.[2] The universe U of $\boldsymbol{P}$ is the set of all data values that can be represented by strings of length at most n (for some given fixed alphabet) and the underlying network consists of all nodes which are represented by strings of length at most n, that is, $\mathcal{N} = \mathbf{dom}_n$. A node κ is responsible for a fact $\boldsymbol{f}$ if $\mathcal{A}_{\boldsymbol{P}}$, on input $(\kappa, \boldsymbol{f})$, has an accepting run of at most $|(\kappa, \boldsymbol{f})|^k$ steps. Clearly, each policy of $\mathcal{P}_{\text{fin}}$ can be described in $\mathcal{P}^2_{\text{npoly}}$. Let $\mathfrak{P}_{\text{npoly}}$ denote the set[3] $\{\mathcal{P}^k_{\text{npoly}} \mid k \geq 2\}$ of distribution policies and by $\mathfrak{P}$ the set $\{\mathcal{P}_{\text{fin}}, \mathcal{P}_{\text{rule}}\} \cup \mathfrak{P}_{\text{npoly}}$.

3.4 Parallel-correctness, soundness, and completeness

In this paper, we mainly consider the one-round evaluation algorithm for a query $\mathcal{Q}$ that first distributes (reshuffles) the data over the computing nodes according to $\boldsymbol{P}$, then evaluates Q in a parallel step at every computing node, and finally outputs all facts that are obtained in this way.[4] As formalised next, the one-round evaluation algorithm is correct (sound, complete) if the query $\mathcal{Q}$ is parallel-correct (parallel-sound, parallel-complete) under $\boldsymbol{P}$.

Definition 3.2. Let $\mathcal{Q}$ be a query, I an instance, and $\boldsymbol{P}$ a distribution policy.

- $\mathcal{Q}$ is *parallel-sound on I under $\boldsymbol{P}$* if $\mathcal{Q}(I) \supseteq [\mathcal{Q}, \boldsymbol{P}](I)$.
- $\mathcal{Q}$ is *parallel-complete on I under $\boldsymbol{P}$* if $\mathcal{Q}(I) \subseteq [\mathcal{Q}, \boldsymbol{P}](I)$; and,
- $\mathcal{Q}$ is *parallel-correct on I under $\boldsymbol{P}$* if $\mathcal{Q}(I) = [\mathcal{Q}, \boldsymbol{P}](I)$, that is, if it is parallel-sound and parallel-complete.

Definition 3.3. A query $\mathcal{Q}$ is *parallel-correct (respectively, parallel-sound and parallel-complete) under distribution policy $\boldsymbol{P} = (U, \mathit{rfacts}_{\boldsymbol{P}})$*, if $\mathcal{Q}$ is parallel-correct (respectively, parallel-sound and parallel-complete) on all instances $I \subseteq \mathit{facts}(\mathcal{D}, U)$.

In [5], parallel-correctness is characterised in terms of minimal valuations as defined next:

[1] We note that it is important that for each class of policies there is a fixed k that bounds the exponent in the test algorithm as otherwise we could not expect a polynomial bound for all policies of that class.

[2] For concreteness, say, a non-deterministic Turing machine.

[3] Since "linear time" is a subtle notion, we rather not consider $\mathcal{P}^1_{\text{npoly}}$.

[4] We note that, since $\boldsymbol{P}$ is defined on the granularity of a fact, the reshuffling does not depend on the current distribution of the data and can be done in parallel as well.

Definition 3.4. Let $\mathcal{Q}$ be a CQ. A valuation V for $\mathcal{Q}$ is *minimal* for $\mathcal{Q}$ if there exists no valuation V' for $\mathcal{Q}$ such that $V(\mathit{head}_\mathcal{Q}) = V'(\mathit{head}_\mathcal{Q})$ and $V'(\mathit{body}_\mathcal{Q}) \subsetneq V(\mathit{body}_\mathcal{Q})$.

The following lemma is key in obtaining the Π^p_2 upper bound on the complexity of testing parallel-correctness for conjunctive queries:

Lemma 3.5 (Characterisation of parallel-correctness for CQs [5])**.** A CQ $\mathcal{Q}$ is parallel-correct under distribution policy $\boldsymbol{P} = (U, \mathit{rfacts}_{\boldsymbol{P}})$ if and only if the following holds:

> For every minimal valuation V for $\mathcal{Q}$ over U, there is a node $\kappa \in \mathcal{N}$ such that $V(\mathit{body}_\mathcal{Q}) \subseteq \mathit{rfacts}_{\boldsymbol{P}}(\kappa)$. (*C1*)

Remark 3.6. Informally, condition (C1) states that there is a node in the network where all facts required for V meet.

3.5 Algorithmic problems

We consider the following decision problems for various sub-classes $\mathcal{C}$ and $\mathcal{C}'$ of $\mathbf{UCQ}^{\neg,\neq}$ and classes $\mathcal{P}$ of distribution policies from $\{\mathcal{P}_{\text{fin}}, \mathcal{P}_{\text{rule}}\} \cup \mathfrak{P}_{\text{npoly}}$.

Containment$(\mathcal{C}, \mathcal{C}')$: **Input:** $\mathcal{Q} \in \mathcal{C}$ and $\mathcal{Q}' \in \mathcal{C}'$ **Question:** Is $\mathcal{Q} \subseteq \mathcal{Q}'$?

Parallel-Sound$(\mathcal{C}, \mathcal{P})$: **Input:** $\mathcal{Q} \in \mathcal{C}$, $\boldsymbol{P} \in \mathcal{P}$ **Question:** Is $\mathcal{Q}$ parallel-sound under $\boldsymbol{P}$?

Parallel-Complete$(\mathcal{C}, \mathcal{P})$: **Input:** $\mathcal{Q} \in \mathcal{C}$, $\boldsymbol{P} \in \mathcal{P}$ **Question:** Is $\mathcal{Q}$ parallel-complete under $\boldsymbol{P}$?

Parallel-Correct$(\mathcal{C}, \mathcal{P})$: **Input:** $\mathcal{Q} \in \mathcal{C}$, $\boldsymbol{P} \in \mathcal{P}$ **Question:** Is $\mathcal{Q}$ parallel-correct under $\boldsymbol{P}$?

4 Parallel-correctness: unions of conjunctive queries

Parallel-correctness of unions of conjunctive queries (without negation) reduces to parallel-completeness for the simple reason that these queries are monotone and therefore parallel-sound for every distribution policy. We show below that parallel-completeness remains in Π^p_2. Hardness already follows from Π^p_2-hardness of Parallel-Correct$(\mathbf{CQ}, \mathcal{P}_{\text{fin}})$ [5].

As a UCQ is parallel-complete under a policy $\boldsymbol{P}$ when all its disjuncts are, it might be tempting to assume that this condition is also necessary. However, as the following example illustrates, this is not the case.

Example 4.1. Let $\mathcal{Q} = \mathcal{Q}_1 \cup \mathcal{Q}_2$, where $\mathcal{Q}_1$ and $\mathcal{Q}_2$ are the following CQs:

$$\begin{array}{lrcl} \mathcal{Q}_1: & H(x,x) & \leftarrow & R(x,x), \\ \mathcal{Q}_2: & H(y,z) & \leftarrow & R(y,z), S(y,z). \end{array}$$

Further, let $\boldsymbol{P}$ be the policy over network $\{\kappa_1, \kappa_2\}$ that maps facts $R(a, a)$ to node κ_1, for all $a \in \mathbf{dom}$, and all other R-facts and all S-facts to node κ_2.

We argue that $\mathcal{Q}$ is parallel-complete under $\boldsymbol{P}$ on all instances. Indeed, assume $H(a, b) \in \mathcal{Q}(I)$ for some instance I and $a, b \in \mathbf{dom}$. If $a \neq b$, only the valuation $\{y \mapsto a, z \mapsto b\}$ can derive $H(a, b)$. This means that $\{R(a, b), S(a, b)\} \subseteq I$. Furthermore, $\{R(a, b), S(a, b)\} \subseteq \mathit{rfacts}_{\boldsymbol{P}}(\kappa_2)$. Hence, $H(a, b) \in \mathcal{Q}(\mathit{loc\text{-}inst}_{\boldsymbol{P},I}(\kappa_2))$. If $a = b$, then $R(a, a) \in I$. So, $R(a, a) \in \mathit{rfacts}_{\boldsymbol{P}}(\kappa_1)$ and $H(a, a) \in \mathcal{Q}(\mathit{loc\text{-}inst}_{\boldsymbol{P},I}(\kappa_1))$. On the other hand, $\mathcal{Q}_2$ is not parallel-complete under $\boldsymbol{P}$ on instance $I = \{R(0, 0), S(0, 0)\}$. Indeed, $H(0, 0) \in \mathcal{Q}_2(I)$ but $\mathcal{Q}_2\big(\mathit{loc\text{-}inst}_{\boldsymbol{P},I}(\kappa_1)\big) = \mathcal{Q}_2\big(\{R(0, 0)\}\big) = \emptyset$ and $\mathcal{Q}_2\big(\mathit{loc\text{-}inst}_{\boldsymbol{P},I}(\kappa_2)\big) = \mathcal{Q}_2\big(\{S(0, 0)\}\big) = \emptyset$. □

We recall from Section 3.2 that disjuncts in unions of conjunctive queries use disjoint variable sets and a valuation for $\mathcal{Q}$ is a valuation for exactly one disjunct. As formalised next, the notion of minimality for valuations given in Definition 3.4 naturally extends to $\mathbf{UCQ}^{\neq}$.

Definition 4.2. Let $\mathcal{Q} = \bigcup_{i=1}^{n} \mathcal{Q}_i$ be a UCQ$^{\neq}$. A valuation V_i for $\mathcal{Q}_i$, with $i \in \{1, \ldots, n\}$, is *minimal* for $\mathcal{Q}$, if for no $j \in \{1, \ldots, n\}$ there is a valuation V_j for $\mathcal{Q}_j$, such that $V_j(\mathit{head}_{\mathcal{Q}_j}) = V_i(\mathit{head}_{\mathcal{Q}_i})$ and $V_j(\mathit{body}_{\mathcal{Q}_j}) \subsetneq V_i(\mathit{body}_{\mathcal{Q}_i})$.

Example 4.3. Consider a simple UCQ$^{\neq}$ $\mathcal{Q} = \mathcal{Q}_1 \cup \mathcal{Q}_2$ where $\mathcal{Q}_1, \mathcal{Q}_2 \in \mathbf{CQ}^{\neq}$ are as follows:

$$\begin{array}{lllll} \mathcal{Q}_1: & H(u, v) & \leftarrow & R(u, v), R(v, u), R(u, u), \\ \mathcal{Q}_2: & H(x, y) & \leftarrow & R(x, y), R(y, z), y \neq z. \end{array}$$

Valuation $V_2 \stackrel{\text{def}}{=} \{x \mapsto 0, y \mapsto 0, z \mapsto 1\}$ is not minimal for $\mathcal{Q}$ because valuation $V_1 \stackrel{\text{def}}{=} \{u \mapsto 0, v \mapsto 0\}$ derives the same fact $H(0, 0)$ requiring only $\{R(0, 0)\} \subsetneq \{R(0, 0), R(0, 1)\}$. Similarly, valuation $W_1 \stackrel{\text{def}}{=} \{u \mapsto 0, v \mapsto 1\}$, requiring $\{R(0, 1), R(1, 0), R(0, 0)\}$, is not minimal for $\mathcal{Q}$ because valuation $W_2 \stackrel{\text{def}}{=} \{x \mapsto 0, y \mapsto 1, z \mapsto 0\}$ only requires $\{R(0, 1), R(1, 0)\}$. □

The notion of minimality leads to basically the same simple characterisation of parallel-completeness:

Lemma 4.4. A UCQ$^{\neq}$ $\mathcal{Q}$ is parallel-correct under distribution policy $\boldsymbol{P} = (U, \mathit{rfacts}_{\boldsymbol{P}})$ if and only if the following holds:

> For every minimal valuation V for $\mathcal{Q}$ over U, there is a node $\kappa \in \mathcal{N}$ such that $V(\mathit{body}_{\mathcal{Q}}) \subseteq \mathit{rfacts}_{\boldsymbol{P}}(\kappa)$. $(C1')$

Proof. (If) Assume $(C1')$ holds. Because of monotonicity, we only need to show that $\mathcal{Q}(I) \subseteq \bigcup_{\kappa \in \mathcal{N}} \mathcal{Q}(\mathit{loc\text{-}inst}_{\boldsymbol{P},I}(\kappa))$ for every instance I. To this end, let $\boldsymbol{f}$ be an arbitrary fact that is derived by some valuation V for $\mathcal{Q}$ on I. Then, there is also a minimal valuation V' that is satisfying on I and which derives $\boldsymbol{f}$. Because of $(C1')$, there is a node $\kappa \in \mathcal{N}$ where all facts required by V' meet (cf. Remark 3.6). Hence, $\boldsymbol{f} \in \bigcup_{\kappa \in \mathcal{N}} \mathcal{Q}(\mathit{loc\text{-}inst}_{\boldsymbol{P},I}(\kappa))$, i.e. query $\mathcal{Q}$ is parallel-correct under policy $\boldsymbol{P}$.

(Only if) For a proof by contraposition, suppose that there is a minimal valuation V' for $\mathcal{Q}$ for which the required facts do *not* meet under $\boldsymbol{P}$. Consider the input instance $I = V'(body_{\mathcal{Q}})$. By definition of minimality, there is no valuation that agrees on the head variables and is satisfying for $\mathcal{Q}$ on a strict subset of $V'(body_{\mathcal{Q}})$. Therefore, $V'(head_{\mathcal{Q}})$ is in $\mathcal{Q}(I)$ but it is not derived on any node and thus query $\mathcal{Q}$ is not parallel-complete under policy $\boldsymbol{P}$. □

The characterisation in Lemma 4.4, in turn, can be used to prove a Π_2^p upper bound.

Lemma 4.5. Parallel-Correct($\mathbf{UCQ}^{\neq}, \mathcal{P}$) is in Π_2^p, for every $\mathcal{P} \in \mathfrak{P}$.

Proof. It suffices to show that the complement of Parallel-Complete($\mathbf{UCQ}^{\neq}, \mathcal{P}^k_{\text{npoly}}$) is in Σ_2^p for arbitrary $k \geq 2$. Let $\boldsymbol{P} = (n, T)$ be a policy from $\mathcal{P}^k_{\text{npoly}}$. We have to consider only instances whose data values can be represented by strings of length n over networks whose nodes can be represented by strings of length n.

By Lemma 4.4, a query $\mathcal{Q}$ is not parallel-correct under distribution policy $\boldsymbol{P}$ if and only if there exists a minimal valuation V that satisfies $\mathcal{Q}$ on some instance I with $adom(I) \subseteq \mathbf{dom}_n$ such that no node in $\mathbf{dom}_n$ is responsible for all facts from $V(body_{\mathcal{Q}})$.

First, the algorithm non-deterministically guesses a valuation V, which can be represented by a string in length polynomial in $\mathcal{Q}$ and n. Subsequently, it checks for all valuations V', all nodes κ, and all strings x of polynomial length whether V' contradicts minimality of V (in which case the algorithm rejects the input) and, by use of algorithm T, whether node κ is not responsible for at least one fact from $V(body_{\mathcal{Q}})$ (if so, the algorithm continues, otherwise it rejects). All tests can be done in polynomial time. □

From [5] we know the following result.

Theorem 4.6 ([5])**.** Parallel-Correct($\mathbf{CQ}, \mathcal{P}_{\text{fin}}$) is Π_2^p-complete.

Together with Lemma 4.5 we get the following result.

Theorem 4.7. Parallel-Correct($\mathbf{UCQ}^{\neq}, \mathcal{P}$) is Π_2^p-complete, for every $\mathcal{P} \in \mathfrak{P}$.

5 Containment of CQ$^{\neg}$ and UCQ$^{\neg}$

In this section, we establish the complexity of containment for $\mathbf{CQ}^{\neg}$ and $\mathbf{UCQ}^{\neg}$. We need these results to establish lower bounds on parallel-correctness and its constituents in the next section. Whereas containment for $\mathbf{CQ}$ has been intensively studied in the literature, the analogous problems for $\mathbf{CQ}^{\neg}$ and $\mathbf{UCQ}^{\neg}$ have hardly been addressed and seem to belong to folklore. In fact, we only found a reference of a complexity result for containment of $\mathbf{CQ}^{\neg}$ in [18], where a Π_2^p-algorithm for the problem is given, based on observations in [15], and the existence of a matching lower bound is mentioned. However, as we show below, although the problem is indeed in Π_2^p for queries defined over a fixed schema (or

when the arity of relations is bounded), it is coNEXPTIME-complete in the general case.

We first show the lower bounds. They actually already hold for Boolean queries. We show that Containment($\mathbf{BCQ}^\neg, \mathbf{UBCQ}^\neg$) is coNEXPTIME-hard by a reduction from the succinct 3-colorability problem and afterwards that Containment($\mathbf{BCQ}^\neg, \mathbf{UBCQ}^\neg$) can be reduced to Containment($\mathbf{BCQ}^\neg, \mathbf{BCQ}^\neg$). Here, $\mathbf{BCQ}^\neg$ and $\mathbf{UBCQ}^\neg$ denote the class of Boolean CQ$^\neg$s and unions of Boolean CQ$^\neg$s, respectively. Together this establishes that Containment($\mathbf{BCQ}^\neg, \mathbf{BCQ}^\neg$) and therefore also Containment($\mathbf{CQ}^\neg, \mathbf{CQ}^\neg$) are coNEXPTIME-hard.

Proposition 5.1. Containment($\mathbf{BCQ}^\neg, \mathbf{UBCQ}^\neg$) is coNEXPTIME-hard.

Proof. The proof is by a reduction from the succinct 3-colorability problem, which asks, whether a graph G, which is implicitly given by a circuit with binary AND- and OR- and unary NEG-gates, is 3-colorable. The latter problem is known to be NEXPTIME-complete [17]. We say that a circuit C, with 2ℓ Boolean inputs, describes a graph $G = (N, E)$, when $N = \{0,1\}^\ell$, and there is an edge $(n_1, n_2) \in N^2$ if and only if C outputs true on input $n_1 n_2$.

Let C be an input for the succinct 3-colorability problem with 2ℓ Boolean inputs. We construct queries $\mathcal{Q}_1$ and $\mathcal{Q}_2$ such that $\mathcal{Q}_1 \not\subseteq \mathcal{Q}_2$ if and only if the graph described by C is 3-colorable.

Both queries are over schema $\mathcal{D}$, which consists of relation names $\texttt{DomainValues}^{(3)}$, $\texttt{Bool}^{(1)}$, $\texttt{And}^{(3)}$, $\texttt{Or}^{(3)}$, $\texttt{Neg}^{(2)}$, and $\texttt{Label}^{(\ell+1)}$. Intuitively, satisfaction of $\mathcal{Q}_1$ will guarantee that there is a tuple (a_0, a_1, a_2) with three different values in relation DomainValues. We will use, for some such tuple, a_0, a_1, a_2 as colors and a_0, a_1 as truth values. We will often assume without loss of generality that $(a_0, a_1, a_2) = (0, 1, 2)$. In particular, for such a tuple, a_0 is interpreted as false while a_1 is interpreted as true. The unary relation Bool will be forced by $\mathcal{Q}_1$ to contain at least a_0 and a_1.

Relations And, Or, and Neg are intended to represent the respective logical functions. The first two attributes represent input values, and the last attribute represents the output. Again, $\mathcal{Q}_1$ will guarantee that at least all triples of Boolean values that are consistent with the semantics of AND, OR, and NEG are present in these relations. Tuples in relation Label represent nodes together with their respective color (one can think of the representation of a node by ℓ-ary addresses over a ternary alphabet).
We define query $\mathcal{Q}_1$ as follows:

$$
\begin{aligned}
T() \leftarrow\ & \texttt{DomainValues}(w_0, w_1, w_2), \neg\texttt{DomainValues}(w_1, w_0, w_2),\\
& \neg\texttt{DomainValues}(w_2, w_1, w_0), \neg\texttt{DomainValues}(w_0, w_2, w_1),\\
& \texttt{Bool}(w_0), \texttt{Bool}(w_1), \texttt{Neg}(w_1, w_0), \texttt{Neg}(w_0, w_1),\\
& \texttt{And}(w_0, w_0, w_0), \texttt{And}(w_0, w_1, w_0), \texttt{And}(w_1, w_0, w_0), \texttt{And}(w_1, w_1, w_1),\\
& \texttt{Or}(w_0, w_0, w_0), \texttt{Or}(w_0, w_1, w_1), \texttt{Or}(w_1, w_0, w_1), \texttt{Or}(w_1, w_1, w_1).
\end{aligned}
$$

It is easy to see that $\mathcal{Q}_1$ enforces the conditions mentioned above.

In the following, we denote sequences $x_1, \ldots, x_\ell$ of ℓ variables by $\boldsymbol{x}$.

We define $\mathcal{Q}_2$ as the union of the queries $\mathcal{Q}_2^1$ and $\mathcal{Q}_2^2$, where subquery $\mathcal{Q}_2^1$ is defined as:

$$T() \leftarrow \texttt{Bool}(x_1), \texttt{Bool}(x_2), \ldots, \texttt{Bool}(x_\ell), \texttt{DomainValues}(y_r, y_g, y_b),$$
$$\neg\texttt{Label}(\boldsymbol{x}, y_r), \neg\texttt{Label}(\boldsymbol{x}, y_g), \neg\texttt{Label}(\boldsymbol{x}, y_b).$$

Intuitively, $\mathcal{Q}_2^1$ can be satisfied in a database if for some node, represented by $\boldsymbol{x}$, there is no color.

Subquery $\mathcal{Q}_2^2$ deals with the correctness of a coloring and uses a set CIRCUIT of atoms that is intended to check whether for two nodes u and v, represented by $\boldsymbol{y}$ and $\boldsymbol{z}$, respectively, there is an edge between u and v.

To this end, CIRCUIT uses the variables $y_1, \ldots, y_\ell$, $z_1, \ldots, z_\ell$, representing the input and, at the same time, the 2ℓ input gates of C, and an additional variable u_i, for each gate of C, with the exception of the output gate. The output gate is represented by variable w_1. For each AND-gate represented by variable v_1 with incoming edges from gates represented by variables u_1 and u_2, CIRCUIT contains an atom $\texttt{And}(u_1, u_2, v_1)$. Likewise for OR- and NEG-gates.

Subquery $\mathcal{Q}_2^2$ is defined as:

$$T() \leftarrow \texttt{DomainValues}(w_0, w_1, w_2), \text{CIRCUIT}, \texttt{Label}(\boldsymbol{y}, u), \texttt{Label}(\boldsymbol{z}, u).$$

Intuitively, $\mathcal{Q}_2^2$ returns true when two nodes, witnessed to be adjacent by the circuit, have the same color.

We show in the appendix that C is 3-colorable if and only if $\mathcal{Q}_1 \not\subseteq \mathcal{Q}_2$. □

Next, we provide the above mentioned reduction.

Proposition 5.2. CONTAINMENT($\mathbf{BCQ}^\neg, \mathbf{UBCQ}^\neg$) $\leq_p$ CONTAINMENT($\mathbf{BCQ}^\neg, \mathbf{BCQ}^\neg$)

Proof. Let $\mathcal{Q}_1$ be in $\mathbf{BCQ}^\neg$ and $\mathcal{Q}_2 = \bigcup_{i=1}^m \mathcal{Q}_2^i$ be in $\mathbf{UBCQ}^\neg$ over some database schema $\mathcal{D}$. Recall our assumption, that each disjunct is defined over a disjoint set of variables. Next, we construct CQs $\mathcal{Q}_1'$ and $\mathcal{Q}_2'$ such that $\mathcal{Q}_1' \subseteq \mathcal{Q}_2'$ if and only if, $\mathcal{Q}_1 \subseteq \mathcal{Q}_2$.

We explain the intuition behind the reduction by means of an example. To this end, let $\mathcal{Q}_1$ be $H() \leftarrow A(x, y)$ and let $\mathcal{Q}_2$ be the $\mathcal{Q}_2^1 \cup \mathcal{Q}_2^2$, where $\mathcal{Q}_2^1$ is $H() \leftarrow A(u_1, v_1), B(u_1, v_1)$ and $\mathcal{Q}_2^2$ is $H() \leftarrow A(u_2, v_2), \neg B(u_2, v_2)$, both formulated over the schema $\mathcal{D} = \{A^{(2)}, B^{(2)}\}$. The query $\mathcal{Q}_2'$ takes the following form:

$$H() \leftarrow \texttt{Active}(x_0, x_1; \ell_1, \ell_2), \underbrace{\alpha(\ell_1, \mathcal{Q}_2^1)}_{\mathcal{Q}_{2,1}'}, \underbrace{\alpha(\ell_2, \mathcal{Q}_2^2)}_{\mathcal{Q}_{2,2}'},$$

where $\alpha(w, \mathcal{Q})$ denotes the modification of the body of $\mathcal{Q}$ by replacing every atom $\texttt{R}(\boldsymbol{x})$ by $\texttt{R}'(w, \boldsymbol{x})$. Both queries are defined over the schema $\mathcal{D}' = \{A'^{(3)}, B'^{(3)}, \texttt{Active}^{(4)}\}$. Notice that $\mathcal{Q}_2'$ contains a concatenation of the disjuncts of $\mathcal{Q}_2$. In addition, relations A and B are extended with a new first column with the purpose of labelling tuples. This labelling allows to encode two (or

even more) instances over $\mathcal{D}$ by one instance over $\mathcal{D}'$. Specifically, $body_{\mathcal{Q}'_1}$ (not shown) is constructed in such a way that when there is a satisfying valuation for $\mathcal{Q}'_1$ there are two different data values, say 0 and 1. So, an instance I over $\mathcal{D}$ can be encoded as $I^0 = \{A'(0,a,b) \mid A(a,b) \in I\} \cup \{B'(0,a,b) \mid B(a,b) \in I\}$ or as $I^1 = \{A'(1,a,b) \mid A(a,b) \in I\} \cup \{B'(1,a,b) \mid B(a,b) \in I\}$. In addition, when there is a satisfying valuation for $\mathcal{Q}'_1$, there is an instance I_2 on which every disjunct of $\mathcal{Q}_2$ is true, and there is an instance I_1 on which $\mathcal{Q}_1$ is true. So, both $\mathcal{Q}'_{2,1}$ and $\mathcal{Q}'_{2,2}$ evaluate to true on I_2^0 when ℓ_1 and ℓ_2 are interpreted by label 0. However, for $\mathcal{Q}_1$ to be contained in $\mathcal{Q}_2$, we need that at least one of the disjuncts $Q'_{2,1}$ or $Q'_{2,2}$ evaluates to true over I_1^1, that is, when its labelling variable is interpreted as 1. Atom $\texttt{Active}(x_0, x_1; \ell_1, \ell_2)$ will ensure that x_0 and x_1 correspond with the values 0 and 1, and that at least one of the labelling variables ℓ_1 or ℓ_2 is equal to 1. In other words, `Active` chooses which disjunct to activate over I_1. So, at least one disjunct of $\mathcal{Q}_2$ evaluates to true on the instance I_1 on which $\mathcal{Q}_1$ is satisfied.

We explain the reduction in more detail in the appendix. □

Combining Propositions 5.1 and 5.2 we get the following corollary:

Corollary 5.3. CONTAINMENT($\mathbf{CQ}^{\neg}, \mathbf{CQ}^{\neg}$) is coNEXPTIME-hard.

The corresponding upper bounds hold also in the presence of inequalities and are shown by small model (i.e., counter-example) properties. To this end, we make use of a restricted monotonicity property of $\text{UCQ}^{\neg,\neq}$s which was already observed in Proposition 2.4 of [3]. For an instance I and a set D of data values we denote by $I_{|D}$ the restriction of I to facts that only use values from D.

Lemma 5.4 ([3])**.** For $\mathcal{Q} \in \mathbf{UCQ}^{\neg,\neq}$, I an instance with a compatible schema, and D a set of data values, it holds that $\mathcal{Q}(I_{|D}) \subseteq \mathcal{Q}(I)$.

Proof. Let $\boldsymbol{f} \in \mathcal{Q}(I_{|D})$ via a valuation V for a disjunct $\mathcal{Q}_i$ of $\mathcal{Q}$. Thus, $V(pos_{\mathcal{Q}_i}) \subseteq I_{|D} \subseteq I$. By definition, every variable x of $\mathcal{Q}_i$ occurs in a positive atom and therefore $V(x) \in D$. Thus, $V(neg_{\mathcal{Q}_i}) \cap I = V(neg_{\mathcal{Q}_i}) \cap I_{|D} = \emptyset$ and $\boldsymbol{f} \in \mathcal{Q}(I)$ as claimed. □

Now we can establish the following small model property for testing containment.

Lemma 5.5. Let $\mathcal{Q}_1, \mathcal{Q}_2 \in \mathbf{UCQ}^{\neg,\neq}$. If there is an instance I, where $\mathcal{Q}_1(I) \not\subseteq \mathcal{Q}_2(I)$, then there is also an instance $J \subseteq I$, where $\mathcal{Q}_1(J) \not\subseteq \mathcal{Q}_2(J)$, and $|adom(J)| \leq varmax(\mathcal{Q}_1)$.

Proof. Let I be as in the lemma and let $\boldsymbol{f}$ be a fact with $\boldsymbol{f} \in \mathcal{Q}_1(I)$ and $\boldsymbol{f} \notin \mathcal{Q}_2(I)$. Let V be a valuation that derives $\boldsymbol{f}$ via some disjunct $\mathcal{Q}_1^i$ of $\mathcal{Q}_1$.

Let $D \stackrel{\text{def}}{=} adom(V(pos_{\mathcal{Q}_1^i}))$ and $J \stackrel{\text{def}}{=} I_{|D}$ the set of all facts in I using only values from $adom(V(pos_{\mathcal{Q}_i}))$. By definition, $|adom(J)| \leq varmax(\mathcal{Q}_1)$. Clearly, V is still a satisfying valuation for $\mathcal{Q}_1^i$ over J. However, by Lemma 5.4, $\boldsymbol{f} \notin \mathcal{Q}_2(J) = \mathcal{Q}_2(I_{|D})$. □

The upper bounds follow easily from Lemma 5.5.

Proposition 5.6. The following upper bounds hold:

1. CONTAINMENT($\mathbf{UCQ}^{\neg,\neq}, \mathbf{UCQ}^{\neg,\neq}$) is in coNEXPTIME.
2. For every k, containment of $\mathbf{UCQ}^{\neg,\neq}$-queries over schemas with arity bound k is in Π_2^p.

Proof. In both cases, we consider the complement of CONTAINMENT($\mathbf{UCQ}^{\neg,\neq}, \mathbf{UCQ}^{\neg,\neq}$). Let $m \stackrel{\text{def}}{=} \mathit{varmax}(\mathcal{Q}_1)$.

1. A NEXPTIME algorithm, on input $\mathcal{Q}_1, \mathcal{Q}_2$, can simply guess an instance J with a domain of at most m elements and a fact $\boldsymbol{f}$, and verifies that $\boldsymbol{f} \in \mathcal{Q}_1(J)$ but $\boldsymbol{f} \notin \mathcal{Q}_2(J)$. For the latter tests, it can simply cycle, in exponential time, to all valuations over J for $\mathcal{Q}_1$ and $\mathcal{Q}_2$.
2. For a fixed arity bound, the minimal counter-example J is of size at most m^k. It can thus be guessed in polynomial time. That $\boldsymbol{f} \in \mathcal{Q}_1(J)$ can be verified non-deterministically. That $\boldsymbol{f} \notin \mathcal{Q}_2(J)$ can be verified by a universal computation in polynomial time.

□

A claim of a Π_2^p upper bound for containment of CQs with negation can be found in [18]. It was not made clear there, that this claim assumes bounded arity of the schema. That the containment problem is Π_2^p-complete for schemas of bounded arity has been explicitly shown in [16]. Clearly, Proposition 5.6.2 follows directly and 5.6.1 is only a variation of it. From Proposition 5.6 and Corollary 5.3 the main result of this section immediately follows.

Theorem 5.7. CONTAINMENT($\mathbf{BCQ}^{\neg}, \mathbf{BCQ}^{\neg}$) and CONTAINMENT($\mathbf{UCQ}^{\neg,\neq}, \mathbf{UCQ}^{\neg,\neq}$) are coNEXPTIME-complete.

Of course, the theorem also holds for all classes $\mathcal{C}$ of queries with $\mathbf{BCQ}^{\neg} \subseteq \mathcal{C} \subseteq \mathbf{UCQ}^{\neg,\neq}$.

6 Parallel-correctness: unions of conjunctive queries with negation

As mentioned in Section 4, for conjunctive queries without negation parallel-soundness always holds and thus parallel-correctness and parallel-completeness coincide, thanks to monotonicity. For queries with negation the situation is different. Distributed evaluation can be complete but not sound, or vice versa. For this reason, we have to distinguish all three problems separately: correctness, soundness, and completeness. However, the complexity is the same in all three cases.

Our results show a second, more crucial difference. Whereas parallel completeness for CQs without negation could be characterised in terms of valuations,

that is, objects of polynomial size, our algorithms for CQs with negation involve counter-examples of exponential size (if the arity of schemas is not bounded) and the coNEXPTIME lower bound results indicate that this is unavoidable. We illustrate the observation that counter-examples might need an exponential number of tuples by the following example.

Example 6.1. Let $\mathcal{Q}$ be the following conjunctive query with negation:

$$H() \leftarrow \texttt{Bool}(w_0, w_0), \texttt{Bool}(w_1, w_1), \texttt{Bool}(x_1, x_1), \ldots, \texttt{Bool}(x_n, x_n), \neg\texttt{Bool}(w_0, w_1), \neg\texttt{Rel}(x_1, \ldots, x_n).$$

Let $\boldsymbol{P}$ be the policy defined over universe $U = \{0, 1\}$ and two-node network $\{\kappa_1, \kappa_2\}$, which distributes all facts except $\texttt{Rel}(0, \ldots, 0)$ to node κ_1 and only fact $\texttt{Rel}(0, \ldots, 0)$ to node κ_2.

Query $\mathcal{Q}$ is not parallel-sound under policy $\boldsymbol{P}$, but the smallest counter-example I is of exponential size as we argue next. Indeed, let $I \stackrel{\text{def}}{=} \{\texttt{Bool}(0, 0), \texttt{Bool}(1, 1)\} \cup \{\texttt{Rel}(a_1, \ldots, a_n) \mid (a_1, \ldots, a_n) \in \{0, 1\}^n\}$. Furthermore, let valuation V map variables w_1 and w_0 to 1 and 0, respectively, and map x_i to 0, for every $i \in \{1, \ldots, n\}$. Then, valuation V satisfies $\mathcal{Q}$ on instance $\mathit{loc\text{-}inst}_{\boldsymbol{P},I}(\kappa_1) = I \setminus \{\texttt{Rel}(0, \ldots, 0)\}$ because neither $\texttt{Bool}(0, 1)$ nor $\texttt{Rel}(0, \ldots, 0)$ is contained in the local instance. Furthermore, there is no satisfying valuation W for $\mathcal{Q}$ on the global instance I because W would have to map each x_i to either 0 or 1 implying that $W\big(\texttt{Rel}(x_1, \ldots, x_n)\big) \in I$.

However, there is no smaller instance: let I^* be some instance over universe U that has a locally satisfying valuation V. The combination of atoms $\texttt{Bool}(w_0, w_0), \texttt{Bool}(w_1, w_1)$, and $\neg\texttt{Bool}(w_0, w_1)$ in query $\mathcal{Q}$ then implies existence of both facts $\texttt{Bool}(0, 0)$ and $\texttt{Bool}(1, 1)$ because variables w_0 and w_1 cannot be mapped onto the same data value.

Assume that fact $\texttt{Rel}(a_1, \ldots, a_n)$, for some $(a_1, \ldots, a_n) \in \{0, 1\}^n$ is missing from I^*. Then the valuation W that maps $w_0 \mapsto 0, w_1 \mapsto 1$ and $x_i \mapsto a_i$, for every $i \in \{1, \ldots, n\}$, satisfies $\mathcal{Q}$ also globally, on instance I^*, and can therefore be no example against parallel-soundness, which contradicts our choice of I^*. Thus, $\texttt{Rel}(a_1, \ldots, a_n) \in I^*$, for every $(a_1, \ldots, a_n) \in \{0, 1\}^n$. We therefore have $I \subseteq I^*$ and, in particular, instance I^* contains at least as many facts as instance I. □

The results of this section are summarised in the following theorem:

Theorem 6.2. For every class $\mathcal{P} \in \{\mathcal{P}_{\text{rule}}\} \cup \mathfrak{P}_{\text{npoly}}$ of distribution policies, the following problems are coNEXPTIME-complete.

- Parallel-Sound($\mathbf{UCQ}^\neg, \mathcal{P}$)
- Parallel-Complete($\mathbf{UCQ}^\neg, \mathcal{P}$)
- Parallel-Correct($\mathbf{UCQ}^\neg, \mathcal{P}$)

Theorem 6.2 follow from Propositions 6.3 and 6.5 below. It also holds for $\mathbf{UCQ}^{\neg,\neq}$. It is easy to show that, when restricted to schemas with some fixed (but sufficiently large, for hardness) arity bound, all these problems are Π_2^p-complete.

6.1 Upper bounds

In this section, we show the upper bounds of Theorem 6.2, summarised in the following proposition.

Proposition 6.3. Parallel-Sound($\mathbf{UCQ}^{\neg,\neq}, \mathcal{P}$), Parallel-Complete($\mathbf{UCQ}^{\neg,\neq}, \mathcal{P}$), and Parallel-Correct($\mathbf{UCQ}^{\neg,\neq}, \mathcal{P}$) are in coNEXPTIME, for every class $\mathcal{P} \in \mathfrak{P}$ of distribution policies. If the arity of schemas is bounded by some fixed number, these problems are in Π_2^p.

Proof. As already indicated above, the proof relies on a bound on the size of a smallest counter-example. More specifically, we first show the following claim.

Claim 6.4. Let $\mathcal{Q} \in \mathbf{UCQ}^{\neg,\neq}$ and let $\boldsymbol{P}$ be an arbitrary distribution policy. Then the following statements hold:

1. If $\mathcal{Q}$ is not parallel-complete under $\boldsymbol{P}$, then there is an instance J over a domain with at most $\mathit{varmax}(\mathcal{Q})$ elements such that $\mathcal{Q}$ is not parallel-complete on J under $\boldsymbol{P}$.
2. If $\mathcal{Q}$ is not parallel-sound under $\boldsymbol{P}$, then there is an instance J over a domain with at most $\mathit{varmax}(\mathcal{Q})$ elements such that $\mathcal{Q}$ is not parallel-sound on J under $\boldsymbol{P}$.

Towards (1) let us assume that $\mathcal{Q}$ is not parallel-complete on some instance I under $\boldsymbol{P}$. Let V be a valuation of a disjunct $\mathcal{Q}_i$ of $\mathcal{Q}$ that derives a fact $\boldsymbol{f}$ globally that is not derived on any node of the network. Let $D \stackrel{\text{def}}{=} \mathit{adom}(V(\mathit{pos}_{\mathcal{Q}_i}))$ and $J \stackrel{\text{def}}{=} I_{|D}$. Clearly, $|D| \leq \mathit{varmax}(\mathcal{Q})$ and V still derives $\boldsymbol{f}$ globally on instance J via $\mathcal{Q}_i$. On the other hand, for every node κ, $\mathcal{Q}\big(\mathit{loc\text{-}inst}_{\boldsymbol{P},J}(\kappa)\big) = \mathcal{Q}\big(\mathit{loc\text{-}inst}_{\boldsymbol{P},I}(\kappa)_{|D}\big) \subseteq \mathcal{Q}\big(\mathit{loc\text{-}inst}_{\boldsymbol{P},I}(\kappa)\big)$, thanks to Lemma 5.4. Therefore $\boldsymbol{f}$ is not derived on κ, and thus J witnesses the lack of parallel-completeness of $\mathcal{Q}$ under $\boldsymbol{P}$.

The proof of (2) is completely analogous. Given a counter-example I and a valuation V that derives a fact $\boldsymbol{f}$ on some node κ via $\mathcal{Q}_i$, for which $\boldsymbol{f}$ is not derived globally, we define $D \stackrel{\text{def}}{=} I_{|\mathit{adom}(V(\mathit{pos}_{\mathcal{Q}_i}))}$ and show that $J \stackrel{\text{def}}{=} I_{|D}$ is the desired counter-example.

In the appendix, we describe an algorithm that tests the complement of parallel completeness non-deterministically. □

6.2 Lower bounds

The lower bounds stated in Theorem 6.2 follow from a polynomial time reduction from problem Containment($\mathbf{BCQ}^{\neg}, \mathbf{BCQ}^{\neg}$), for which we showed coNEXPTIME-hardness in Section 5.

Proposition 6.5. Parallel-Complete($\mathbf{CQ}^{\neg}, \mathcal{P}_{\text{rule}}$), Parallel-Sound($\mathbf{CQ}^{\neg}, \mathcal{P}_{\text{rule}}$), and Parallel-Correct($\mathbf{CQ}^{\neg}, \mathcal{P}_{\text{rule}}$) are coNEXPTIME-hard.

Proof. Interestingly, all three results are shown by the *same* reduction from decision problem Containment($\mathbf{BCQ}^\neg, \mathbf{BCQ}^\neg$).

The basic idea for this reduction is very simple: it combines both queries $\mathcal{Q}_1, \mathcal{Q}_2 \in \mathbf{BCQ}^\neg$ of the given containment instance into a single query $\mathcal{Q} \in \mathbf{BCQ}^\neg$ and infers an appropriate distribution policy $\boldsymbol{P}$. To emulate separate derivation for both queries in the combined query, an activation mechanism is used that resembles the proof of Proposition 5.2. In this fashion, the two queries can be evaluated over different subsets of the considered instance by annotating both the facts in the instance as well as the atoms of the query.

We next describe the reduction in detail. Let thus $\mathcal{Q}_1, \mathcal{Q}_2 \in \mathbf{BCQ}^\neg$ be queries over some schema $\mathcal{D}$ and let $m \stackrel{\text{def}}{=} \max\{\mathit{varmax}(\mathcal{Q}_1), \mathit{varmax}(\mathcal{Q}_2)\}$. Without loss of generality, we assume the variable sets of $\mathcal{Q}_1$ and $\mathcal{Q}_2$ to be disjoint. We will also assume in the following that both $\mathcal{Q}_1$ and $\mathcal{Q}_2$ are satisfiable. This is the case (for $\mathcal{Q}_1$) if and only if $\mathit{pos}_{\mathcal{Q}_1} \cap \mathit{neg}_{\mathcal{Q}_1} = \emptyset$ and can therefore be easily tested in polynomial time. If one of the test fails, some appropriate constant instance of Parallel-Complete($\mathbf{CQ}^\neg, \mathcal{P}_{\text{rule}}$) or one of the other problem variants, respectively, can be computed.

We define a (Boolean) query $\mathcal{Q} \in \mathbf{BCQ}^\neg$ and a policy $\boldsymbol{P} \in \mathcal{P}_{\text{rule}}$ over domain $\{1, \ldots, m\}$ that can be computed from $\mathcal{Q}_1$ and $\mathcal{Q}_2$ in polynomial time. The schema for $\mathcal{Q}$ is $\mathcal{D}' \stackrel{\text{def}}{=} \{R'^{(k+1)} \mid R^{(k)} \in \mathcal{D}\}$. That is, each relation name R of $\mathcal{D}$ occurs as R' in $\mathcal{D}'$ with an arity incremented by one. Additionally, $\mathcal{Q}$ uses relation names $\texttt{Type}$, $\texttt{Start}_1$, $\texttt{Start}_2$, and $\texttt{Stop}$, which we assume not to occur in schema $\mathcal{D}$. Besides the variables of $\mathcal{Q}_1$ and $\mathcal{Q}_2$, query $\mathcal{Q}$ uses variables ℓ_1, ℓ_2, t.

We use the function α, defined in the proof of Proposition 5.2, which adds its first parameter as first component to every tuple in its second parameter and translates relation names R into R'. In Proposition 5.2, the first parameter was always a variable and the second a set of atoms, but we use α also for a data value as first and a set of facts as second parameter in the obvious way. We write α_a^{-1} for the function mapping sets of facts over $\mathcal{D}'$ to sets of facts over $\mathcal{D}$, by selecting, from a set of facts, all facts with first parameter a, deleting this parameter and replacing each name R' by R. Finally, $\pi_a(I) \stackrel{\text{def}}{=} \alpha\big(a, \alpha_a^{-1}(I)\big)$ is the restriction of I to all facts with a in their first component.

The combined query $\mathcal{Q}$ has $\mathit{head}_{\mathcal{Q}} \stackrel{\text{def}}{=} H()$ and body

$$\begin{array}{rcl} \mathit{body}_{\mathcal{Q}} & \stackrel{\text{def}}{=} & \alpha(\ell_1, \mathit{body}_{\mathcal{Q}_1}) \cup \alpha(\ell_2, \mathit{body}_{\mathcal{Q}_2}) \\ & \cup & \underbrace{\{\texttt{Type}(t), \texttt{Start}_1(\ell_1), \texttt{Start}_2(\ell_1), \texttt{Start}_2(\ell_2)\}}_{\mathcal{A}} \cup \underbrace{\{\neg\texttt{Stop}(\ell_1), \neg\texttt{Stop}(\ell_2)\}}_{\mathcal{A}^\neg}. \end{array}$$

Policy $\boldsymbol{P}$ is defined over universe $U \stackrel{\text{def}}{=} \{1, \ldots, m\}$, schema $\mathcal{D}' \cup \{\texttt{Type}, \texttt{Start}_1, \texttt{Start}_2, \texttt{Stop}\}$ and network $\mathcal{N} \stackrel{\text{def}}{=} \{\kappa_1, \ldots, \kappa_m, \sigma_1, \ldots, \sigma_m, \rho\}$. Facts are distributed as follows:

- Every node κ_i is responsible for the facts $\texttt{Type}(1), \texttt{Start}_1(i), \texttt{Start}_2(i), \texttt{Stop}(i)$, and all facts from $\mathit{facts}(\mathcal{D}', U)$.
- Every node σ_i is responsible for the facts $\texttt{Type}(2), \texttt{Start}_1(i), \texttt{Stop}(i)$, all $\texttt{Start}_2$-facts, and all facts from $\mathit{facts}(\mathcal{D}', U)$.

- Finally, node ρ is responsible for facts $\texttt{Type}(3), \ldots, \texttt{Type}(m)$, and *all* facts over other relation names.

It is easy to see that $\boldsymbol{P}$ can be expressed by a polynomial number of rules and that $\mathcal{Q}$ and $\boldsymbol{P}$ can be computed in polynomial time. In the appendix, we show that the described function is indeed the desired reduction. □

7 Full conjunctive queries

In this section, we focus attention on full conjunctive queries, in an attempt to lower the complexity of testing parallel-correctness. Requiring queries to be full is a very natural restriction which is known to have practical benefits. For example, the Hypercube algorithm, which describes an optimal way to compute CQs in a setting very similar to ours, completely ignores projections when shuffling data, and only applies them when computing the query locally. The latter is possible because correctness for the full-variant of a query is in a sense more strict than correctness for the query itself.

Formally, a (union of) conjunctive queries is called *full* if all variables of the body also occur in the head. We denote by $\mathbf{FCQ}^{\neg,\neq}$ and $\mathbf{UFCQ}^{\neg,\neq}$ the class of full $\mathrm{CQ}^{\neg,\neq}$ and full $\mathrm{UCQ}^{\neg,\neq}$ queries, respectively, and likewise for other fragments.

The presentation is similar to that of Section 5 and 6. First, we establish the complexity of query containment. Then, we show that containment reduces to parallel-correctness (and variants). Finally, we obtain matching upper bounds.

The following theorem shows that unlike for general conjunctive queries the complexity of deciding containment for $\mathrm{FCQ}^{\neg}$ and $\mathrm{UFCQ}^{\neg}$ do not coincide.

Theorem 7.1. 1. Containment($\mathbf{FCQ}^{\neg}, \mathbf{FCQ}^{\neg}$) is in P;

2. Containment($\mathbf{FCQ}^{\neg}, \mathbf{UFCQ}^{\neg}$) is coNP-complete; and

3. Containment($\mathbf{UFCQ}^{\neg}, \mathbf{UFCQ}^{\neg}$) is coNP-complete.

All these results also hold for queries with inequalities.

As one can reduce from Containment($\mathbf{FCQ}^{\neg}, \mathbf{UFCQ}^{\neg}$) to parallel-soundness, completeness, and correctness, we obtain the following hardness results:

Proposition 7.2. Parallel-Sound($\mathbf{UFCQ}^{\neg}, \mathcal{P}$), Parallel-Complete($\mathbf{UFCQ}^{\neg}, \mathcal{P}$), and Parallel-Correct($\mathbf{UFCQ}^{\neg}, \mathcal{P}$) are coNP-hard, for every $\mathcal{P} \in \{\mathcal{P}_{\text{rule}}\} \cup \mathfrak{P}_{\text{npoly}}$.

The following theorem determines the complexity for the upper bounds:

Theorem 7.3. The following problems are coNP-complete:

1. Parallel-Sound($\mathbf{UFCQ}^{\neg}, \mathcal{P}_{\text{rule}}$);

2. Parallel-Complete($\mathbf{UFCQ}^{\neg}, \mathcal{P}_{\text{rule}}$);

3. Parallel-Correct($\mathbf{UFCQ}^{\neg}, \mathcal{P}_{\text{rule}}$).

The result also holds for queries with inequalities.

8 Discussion

In this paper, we continued the study of parallel-correctness initiated by Ameloot et al. [5] as a framework for reasoning about one-round evaluation algorithms for conjunctive queries under arbitrary distribution policies. Specifically, we considered the case with union and negation. While parallel-correctness for unions of conjunctive queries can be tested by examining properties of single valuations, just like in the union-free case, the latter no longer holds true when negation is present. Consequently, we obtained that deciding parallel-correctness for unions of conjunctive queries remains in Π_2^p, while the analog problem in the presence of negation is hard for coNEXPTIME. Since conjunctive queries with negation are no longer monotone, we considered the related problems of parallel-completeness and parallel-soundness as well and obtained the same bounds. Interestingly, when negation is present, containment of conjunctive queries can be reduced to parallel-correctness (and its variants) allowing the transfer of lower bounds. We prove that containment for conjunctive queries with negation is hard for coNEXPTIME, which, to the best of our knowledge, is a novel result. In an attempt to lower complexity, we show that parallel-correctness for unions of full conjunctive queries with negation is coNP-complete.

There are quite a number of directions towards future work. While parallel-correctness for first-order logic is undecidable, it would be interesting to determine the exact frontier for decidability. As the considered problem is a static analysis problem that relates to the size of the queries and not to the size of the instances (at least in the setting of $\mathcal{P}_{\text{rule}}$), exponential lower bounds do not necessarily exclude practical application. It could still be interesting to identify settings that would make parallel-correctness tractable. Possibly independent of tractability considerations, such settings could incorporate bag semantics, integrity constraints, or specific classes (and representations) of distribution policies. We also plan to consider evaluation algorithms that use knowledge about the distribution policy to compute better query results, locally. Another direction for future work is to investigate transferability of parallel-correctness for conjunctive queries as defined in [5] in the presence of union and negation.

References

[1] Spark. http://spark.apache.org.

[2] F. N. Afrati and J. D. Ullman. Optimizing joins in a map-reduce environment. In *EDBT*, pages 99–110, 2010.

[3] Foto N. Afrati, Stavros S. Cosmadakis, and Mihalis Yannakakis. On datalog vs. polynomial time. *J. Comput. Syst. Sci.*, 51(2):177–196, 1995.

[4] Foto N. Afrati, Paraschos Koutris, Dan Suciu, and Jeffrey D. Ullman. Parallel skyline queries. In *International Conference on Database Theory (ICDT 2012)*, pages 274–284, 2012.

[5] Tom J. Ameloot, Gaetano Geck, Bas Ketsman, Frank Neven, and Thomas Schwentick. Parallel-correctness and transferability for conjunctive queries. In *Principles of Database Systems (PODS 2015)*, pages 47–58. ACM, 2015.

[6] Tom J. Ameloot, Bas Ketsman, Frank Neven, and Daniel Zinn. Weaker forms of monotonicity for declarative networking: a more fine-grained answer to the CALM-conjecture. In *Principles of Database Systems (PODS 2014)*, pages 64–75, 2014.

[7] Tom J. Ameloot, Frank Neven, and Jan Van den Bussche. Relational transducers for declarative networking. *J. ACM*, 60(2):15, 2013.

[8] Paul Beame, Paraschos Koutris, and Dan Suciu. Communication steps for parallel query processing. In *PODS*, pages 273–284, 2013.

[9] Paul Beame, Paraschos Koutris, and Dan Suciu. Skew in parallel query processing. In *Principles of Database Systems (PODS 2014)*, pages 212–223, 2014.

[10] Ashok K. Chandra and Philip M. Merlin. Optimal implementation of conjunctive queries in relational data bases. In *STOC*, pages 77–90, 1977.

[11] Shumo Chu, Magdalena Balazinska, and Dan Suciu. From theory to practice: Efficient join query evaluation in a parallel database system. In *SIGMOD*, pages 63–78, 2015.

[12] Carles Farré, Werner Nutt, Ernest Teniente, and Toni Urpí. Containment of conjunctive queries over databases with null values. In *11th International Conference on Database Theory, ICDT'07*, pages 389–403, 2007.

[13] Sumit Ganguly, Abraham Silberschatz, and Shalom Tsur. Parallel bottom-up processing of datalog queries. *J. Log. Program.*, 14(1&2):101–126, 1992.

[14] Paraschos Koutris and Dan Suciu. Parallel evaluation of conjunctive queries. In *PODS*, pages 223–234, 2011.

[15] Alon Y. Levy and Yehoshua Sagiv. Queries independent of updates. In *International Conference on Very Large Data Bases (VLDB 1993)*, pages 171–181, 1993.

[16] Marie-Laure Mugnier, Geneviève Simonet, and Michaël Thomazo. On the complexity of entailment in existential conjunctive first-order logic with atomic negation. *Inf. Comput.*, 215:8–31, 2012.

[17] Christos H. Papadimitriou and Mihalis Yannakakis. A note on succinct representations of graphs. *Information and Control*, 71(3):181–185, 1986.

[18] Jeffrey D. Ullman. Information integration using logical views. *Theoretical Computer Science*, 239(2):189–210, 2000.

[19] Fang Wei and Georg Lausen. Containment of conjunctive queries with safe negation. In *Database Theory - ICDT 2003, 9th International Conference, Siena, Italy, January 8-10, 2003, Proceedings*, pages 343–357, 2003.

[20] R. Xin, J. Rosen, M. Zaharia, M. Franklin, S. Shenker, and I. Stoica. Shark: SQL and rich analytics at scale. In *SIGMOD*, 2013.

[21] Daniel Zinn, Todd J. Green, and Bertram Ludäscher. Win-move is cordination-free (sometimes). In *International Conference on Database Theory (ICDT 2012)*, pages 99–113, 2012.

Appendix

9 Proofs for Section 5: Containment of $\mathrm{CQ}^\neg$ and $\mathrm{UCQ}^\neg$

Proof of Proposition 5.1 (continued). It remains to show that the graph represented by circuit C is 3-colorable if and only if $\mathcal{Q}_1 \not\subseteq \mathcal{Q}_2$.

(If) Suppose $\mathcal{Q}_1 \not\subseteq \mathcal{Q}_2$. Thus, for some instance I, $T() \in \mathcal{Q}_1(I)$ while $T() \notin \mathcal{Q}_2(I)$. The former implies that I contains a fact $\texttt{DomainValues}(a_0, a_1, a_2)$, where a_0, a_1, and a_2 are distinct values from **dom**. Without loss of generality, we assume $a_0 = 0, a_1 = 1, a_2 = 2$. The instance I further contains the facts $\texttt{Bool}(0)$ and $\texttt{Bool}(1)$, and all "logical facts" induced by $\mathcal{Q}_1$.

For every vector $\boldsymbol{n} \in \{0,1\}^\ell$, there must be some $c \in \{0,1,2\}$ such that $\texttt{Label}(\boldsymbol{n}; c) \in I$, since otherwise $T() \in \mathcal{Q}_2^1$. Let $\mathit{label} : \{0,1\}^\ell \to \{0,1,2\}$ be chosen such that, for every $\boldsymbol{n} \in \{0,1\}^\ell$, $\mathit{label}(\boldsymbol{n}) = c$, for some $c \in \{0,1,2\}$, for which $\texttt{Label}(\boldsymbol{n}; c) \in I$.

We claim that *label* is a valid coloring of the graph represented by C. Towards a contradiction let us assume that there are two nodes $\boldsymbol{n}$ and $\boldsymbol{n}'$ that are connected by an edge and for which $\mathit{label}(\boldsymbol{n}) = \mathit{label}(\boldsymbol{n}') = c$, for some c. Then $\mathcal{Q}_2^2$ could be satisfied over I by choosing a valuation that corresponds to a computation of C that witnesses that there is an edge between $\boldsymbol{n}$ and $\boldsymbol{n}'$ and mapping u to c, the desired contradiction.

(Only if) Let, for some ℓ, C be a circuit with input length 2ℓ, that describes a 3-colorable graph G. Let $\mathit{label} : \{0,1\}^\ell \to \{0,1,2\}$ be a valid coloring for G. Let I be the database with the following facts:

$$\{\texttt{Label}\big(\boldsymbol{n}; \mathit{label}(\boldsymbol{n})\big) \mid \boldsymbol{n} \in \{0,1\}^\ell\} \cup \{\texttt{DomainValues}(0,1,2), \texttt{Bool}(0), \texttt{Bool}(1)\} \cup$$
$$\{\texttt{Neg}(1,0), \texttt{Neg}(0,1), \texttt{And}(0,0,0), \texttt{And}(0,1,0), \texttt{And}(1,0,0), \texttt{And}(1,1,1)\} \cup$$
$$\{\texttt{Or}(0,0,0), \texttt{Or}(0,1,1), \texttt{Or}(1,0,1), \texttt{Or}(1,1,1)\}.$$

Obviously, $T() \in \mathcal{Q}_1(I)$ and $T() \notin \mathcal{Q}_2^1(I)$. However, since I only contains the "correct" logical facts, to satisfy $\mathcal{Q}_2^2$ it would be necessary to find two nodes with the same label whose adjacency is witnessed by the canonical valuation corresponding to the semantics of C, which does not exist. Thus, $\mathcal{Q}_1 \not\subseteq \mathcal{Q}_2$. □

Proof of Proposition 5.2 (continued). We now explain the reduction in more detail. We assume that CQ$^\neg$ $\mathcal{Q}_2^i$ is satisfiable, for every $i \in \{1, \ldots, m\}$.[5] Set $\mathcal{D}' \stackrel{\text{def}}{=} \{\mathtt{R}'^{(k+1)} \mid \mathtt{R}^{(k)} \in \mathcal{D}\} \cup \{\mathtt{Active}^{(m+2)}\}$. As explained above, the relation `Active` serves as an index for the disjuncts in $\mathcal{Q}_2$. Whereas the first two positions encode the bits zero (0) and one (1), an atom of the form $\mathtt{Active}(0, 1; 0, \ldots, 0, 1, 0, \ldots, 0)$ with a 1 occurring on position $i + 2$ is meant to indicate disjunct i. Recall that $\alpha(w, \mathcal{Q})$ denotes the modification of the body of $\mathcal{Q}$ by replacing every atom $\mathtt{R}(\boldsymbol{x})$ by $\mathtt{R}'(w, \boldsymbol{x})$, while retaining existent negation symbols. We further define[6] mapping α_a^{-1}, for $a \in \mathbf{dom}$, to map sets of facts over $\mathcal{D}'$ to sets of facts over $\mathcal{D}$, by selecting all facts with first parameter a, deleting this parameter and replacing each relation name R' by R.
Now, define $\mathcal{Q}_1'$ as:

$$\begin{aligned} H() \leftarrow & \mathtt{Active}(w_0, w_1; w_1, w_0, \ldots, w_0), \mathtt{Active}(w_0, w_1; w_0, w_1, \ldots, w_0), \ldots, \\ & \mathtt{Active}(w_0, w_1; w_0, w_0, \ldots, w_1, w_0), \mathtt{Active}(w_0, w_1; w_0, \ldots, w_0, w_1), \\ & \neg\mathtt{Active}(w_1, w_0; w_1, w_0, \ldots, w_0), \\ & \alpha(w_1, \mathcal{Q}_1), \alpha(w_0, \mathcal{Q}_2^1), \ldots, \alpha(w_0, \mathcal{Q}_2^m); \end{aligned}$$

and $\mathcal{Q}_2'$ as:

$$H() \leftarrow \mathtt{Active}(w_0, w_1; z_1, \ldots, z_m), \alpha(z_1, \mathcal{Q}_2^1), \ldots, \alpha(z_m, \mathcal{Q}_2^m),$$

where $w_0, w_1, z_1, \ldots, z_m$ are fresh and distinct variables.

In the remainder we argue that $\mathcal{Q}_1' \subseteq \mathcal{Q}_2'$ if and only if $\mathcal{Q}_1 \subseteq \mathcal{Q}_2$:

(If) Suppose $\mathcal{Q}_1 \subseteq \mathcal{Q}_2$. Let I' be an arbitrary instance over $\mathcal{D}'$ for which $H() \in \mathcal{Q}_1'(I)$. Let V_1' be a satisfying valuation for $\mathcal{Q}_1'$ over I'. Since V_1' has to satisfy both literals $\mathtt{Active}(w_0, w_1; w_1, w_0, \ldots, w_0)$ and $\neg\mathtt{Active}(w_1, w_0; w_1, w_0, \ldots, w_0)$, it follows that $V_1'(w_1) \neq V_1'(w_0)$. For simplicity, we assume that $V_1'(w_1) = 1$ and $V_1'(w_0) = 0$.

Thus, I' contains the facts $\mathtt{Active}(0, 1; 1, 0, \ldots, 0), \ldots, \mathtt{Active}(0, 1; 0, \ldots, 0, 1)$, but not $\mathtt{Active}(1, 0; 1, 0, \ldots, 0)$. Furthermore, for every i, $V_1'\big(\alpha(w_0, \mathcal{Q}_2^i)\big)$ consists of facts from I' labelled with 0, and for $V_1'\big(\alpha(w_1, \mathcal{Q}_1)\big)$ consists of facts from I' labelled with 1.

Now, let I be the set $\alpha_1^{-1}(I') = \{R(\mathbf{t}) \mid R'(1, \mathbf{t}) \in I'\}$ consisting of all facts in I' labelled with 1. By construction and since I' contains all facts from $V_1'\big(\alpha(w_1, \mathcal{Q}_1)\big)$, I contains all facts from $V_1'(\mathcal{Q}_1)$. Note that V_1' is a valuation for $\mathcal{Q}_1$ as well. In fact, V_1' is a satisfying valuation for $\mathcal{Q}_1$ over I. Consequently, thanks to $\mathcal{Q}_1 \subseteq \mathcal{Q}_2$, there must be a satisfying valuation V_2 for some disjunct $\mathcal{Q}_2^i$ of $\mathcal{Q}_2$ over I.

We define a satisfying valuation V_2' for $\mathcal{Q}_2'$ over I' witnessing $\mathcal{Q}_1' \subseteq \mathcal{Q}_2'$ as

[5] Notice that a CQ$^\neg$ $\mathcal{Q}$ is satisfiable if and only if $pos_\mathcal{Q} \cap neg_\mathcal{Q} = \emptyset$, which can easily be verified in polynomial time.

[6] This is the same definition as in Section 6.2 and is stated here for convenience only.

follows:

$$V_2'(x) = \begin{cases} 0 & \text{if } x = w_0, \\ 1 & \text{if } x = w_1, \\ 1 & \text{if } x = z_i, \\ 0 & \text{if } x = z_j \text{ and } j \neq i, \\ V_2(x) & \text{if } x \text{ occurs in } \mathcal{Q}_2^i, \\ V_1'(x) & \text{if } x \text{ occurs in } \mathcal{Q}_2^j \text{ and } j \neq i. \end{cases}$$

It is easy to verify that V_2' satisfies $\mathcal{Q}_2'$: the first atom and all conjuncts $\alpha(z_j, \mathcal{Q}_2^j)$ with $j \neq i$ become true since the respective facts were guaranteed by $\mathcal{Q}_1'$. Finally, $\alpha(z_i, \mathcal{Q}_2^i)$ becomes true since V_2 satisfies $\mathcal{Q}_2^i$.

(Only if) The proof is by contraposition, that is, we show that $\mathcal{Q}_1 \not\subseteq \mathcal{Q}_2$ implies $\mathcal{Q}_1' \not\subseteq \mathcal{Q}_2'$. Therefore, let I be an instance, and V_1 a valuation such that V_1 satisfies $\mathcal{Q}_1$ over I, but no $\mathcal{Q}_2^i$ has a satisfying valuation over I. Since every query $\mathcal{Q}_2^i$ is satisfiable, there is, for every i, a satisfying valuation V^i. In fact, these valuations can be chosen with pairwise disjoint range. Now, we define I' as the following set of facts:

$$\begin{array}{rcl} I' & = & \alpha(1, I) \cup \alpha\big(0, V^1(pos_{\mathcal{Q}_2^1})\big) \cup \cdots \cup \alpha\big(0, V^m(pos_{\mathcal{Q}_2^m})\big) \\ & \cup & \{\texttt{Active}(0,1;1,0,\ldots,0), \ldots, \texttt{Active}(0,1;0,\ldots,0,1)\}. \end{array}$$

It is easy to check[7] that from V_1 and the V^i a satisfying valuation for $\mathcal{Q}_1'$ over I' can be constructed. On the other hand, any satisfying valuation of $\mathcal{Q}_2'$ over I' would require to use facts from $\alpha(1, I)$ for at least one $\alpha(z_i, \mathcal{Q}_2^i)$ and would thus induce a valuation of $\mathcal{Q}_2^i$ over I, the desired contradiction. □

10 Proofs for Section 6: Parallel-correctness: unions of conjunctive queries with negation

Proof of Proposition 6.3 (continued). It only remains to describe the algorithm that tests the complement of parallel completeness non-deterministically. On input $\mathcal{Q}$ and $\boldsymbol{P}$ (specified by $(n, T) \in \mathcal{P}^k_{\text{npoly}}$), the algorithm simply guesses an instance J over a domain with at most $varmax(\mathcal{Q})$ values from $\mathbf{dom}_n$, and verifies that J is a counter-example showing that $\mathcal{Q}$ is not parallel-complete under $\boldsymbol{P}$. From Claim 6.4 it follows that this algorithm is correct, since a counter-example must exist if $\mathcal{Q}$ is not parallel-complete under $\boldsymbol{P}$, and the actual data values do not matter. It remains to show the complexity bounds and, in particular, to describe how the verification part can done.

In the general case, without a bound on the arity of the schema, the verification is done as follows. The algorithm guesses a valuation V that produces some fact $\boldsymbol{f}$ globally and which is not derived at any node. To test that $\boldsymbol{f}$ is not derived at any node, the algorithm cycles through all nodes and

[7] This can be proven along the lines of Claim 10.1.

all valuations V over $\mathbf{dom}_n$. The number of combinations is bounded[8] by $2^n \times (2^n)^{varmax(\mathcal{Q})} = 2^{n(varmax(\mathcal{Q})+1)}$. Each test can be performed by a simulation of all runs of T which amounts to at most 2^{n^k} simulations of at most n^k steps each. Altogether the algorithm needs time at most $2^{|(\mathcal{Q},\boldsymbol{P})|^{k+2}}$.

If there is a fixed bound ℓ on the arity of the underlying schema then the maximum size of the minimal counter-example becomes polynomial and the test that $\boldsymbol{f}$ is derived globally can be done non-deterministically in polynomial time and the test that it is not derived locally can be done universally in polynomial time, thus altogether yielding a Π_2^p-computation.

The case of parallel-soundness is completely analogous (using the second statement of Claim 6.4 and the case of parallel-correctness follows since it suffices to test parallel completeness and soundness). □

Proof of Proposition 6.5 (continued). It remains to show that the described function is a reduction from Containment($\mathbf{BCQ}^\neg, \mathbf{BCQ}^\neg$) to all of Parallel-Complete($\mathbf{CQ}^\neg, \mathcal{P}_{\text{rule}}$), Parallel-Sound($\mathbf{CQ}^\neg, \mathcal{P}_{\text{rule}}$), and Parallel-Correct($\mathbf{CQ}^\neg, \mathcal{P}_{\text{rule}}$).

To this end, we show first that containment $\mathcal{Q}_1 \subseteq \mathcal{Q}_2$ implies parallel-completeness and parallel-soundness of query $\mathcal{Q}$ under policy $\boldsymbol{P}$, and that lack of containment, $\mathcal{Q}_1 \not\subseteq \mathcal{Q}_2$, implies that query $\mathcal{Q}$ is neither parallel-complete nor parallel-sound under policy $\boldsymbol{P}$.

In both directions, we will make use of the following easy observations.

Claim 10.1. Let I be an arbitrary instance and $i \in \{1, 2\}$.

1. Let V be a valuation for $\mathcal{Q}$, let $a \stackrel{\text{def}}{=} V(\ell_i)$ and V_i be the restriction of valuation V to variables in $\mathcal{Q}_i$. If V satisfies $\mathcal{Q}$ on I, then V_i satisfies $\mathcal{Q}_i$ on $\alpha_a^{-1}(I)$.

2. Let V_1, V_2 be valuations for queries $\mathcal{Q}_1, \mathcal{Q}_2$. Let a, b be data values such that V_1 and V_2 satisfy $\mathcal{Q}_1$ and $\mathcal{Q}_2$, respectively, on $\alpha_a^{-1}(I)$, and such that $\texttt{Type}(b), \texttt{Start}_1(a), \texttt{Start}_2(a) \in I$, and $\texttt{Stop}(a) \notin I$. Then the valuation W_a^b, that agrees with V_1 and V_2 on all variables in $\mathcal{Q}_1$ and $\mathcal{Q}_2$, respectively, and maps $\ell_1, \ell_2 \mapsto a$, and $t \mapsto b$, satisfies $\mathcal{Q}$ on I.

Let us now assume $\mathcal{Q}_1 \subseteq \mathcal{Q}_2$. To show that $\mathcal{Q}$ is parallel-complete under $\boldsymbol{P}$, let V be a valuation that globally satisfies $\mathcal{Q}$ on some arbitrary instance I over U. Let $a \stackrel{\text{def}}{=} V(\ell_1)$ and $b \stackrel{\text{def}}{=} V(t)$. Satisfaction of $\mathcal{Q}$ on I by V then particularly implies $\texttt{Type}(b), \texttt{Start}_1(a), \texttt{Start}_2(a) \in I$ and $\texttt{Stop}(a) \notin I$. Let V_1 be the satisfying valuation for $\mathcal{Q}_1$ on instance $\alpha_a^{-1}(I)$, as given by Claim 10.1.1. Since $\mathcal{Q}_1 \subseteq \mathcal{Q}_2$, there exists a valuation V_2 that satisfies $\mathcal{Q}_2$ on instance $\alpha_a^{-1}(I)$. By Claim 10.1.2, valuation $W \stackrel{\text{def}}{=} W_a^b$ satisfies $\mathcal{Q}$ on $\pi_a(I) \cup \{\texttt{Type}(b), \texttt{Start}_1(a), \texttt{Start}_2(a)\}$. If $W(t) = 1$, then node κ_a is responsible for these facts; if $W(t) = 2$, then node σ_a is responsible for these facts; and otherwise node ρ is responsible for these facts. Hence, query $\mathcal{Q}$ is parallel-complete under policy $\boldsymbol{P}$.

The proof that $\mathcal{Q}$ is parallel-sound under $\boldsymbol{P}$ is similar. To this end, let I be a global instance and V be a valuation that satisfies $\mathcal{Q}$ on the local instance I_κ

[8] We assume a binary alphabet, here.

of some node $\kappa \in \mathcal{N}$. Let $a \stackrel{\text{def}}{=} V(\ell_1)$ and $b \stackrel{\text{def}}{=} V(t)$. By definition of $\boldsymbol{P}$, we can infer $\texttt{Type}(b), \texttt{Start}_1(a), \texttt{Start}_2(a) \in I$ and $\texttt{Stop}(a) \notin I$ from satisfaction of $\mathcal{Q}$ by V. Let V_1 be the satisfying valuation for $\mathcal{Q}_1$ on instance $\alpha_a^{-1}(I_\kappa)$, as given by Claim 10.1.1. Since $\mathcal{Q}_1 \subseteq \mathcal{Q}_2$, there exists a valuation V_2 that satisfies $\mathcal{Q}_2$ on instance $\alpha_a^{-1}(I_\kappa)$. By definition of $\boldsymbol{P}$, we have $\alpha_a^{-1}(I_\kappa) = \alpha_a^{-1}(I)$, and therefore V_2 satisfies $\mathcal{Q}_2$ on $\alpha_a^{-1}(I)$. Thus, by Claim 10.1.2, valuation $W \stackrel{\text{def}}{=} W_a^b$ satisfies $\mathcal{Q}$ on I. Hence, query $\mathcal{Q}$ is parallel-sound under policy $\boldsymbol{P}$.

Let now $\mathcal{Q}_1 \not\subseteq \mathcal{Q}_2$ and let I_1 be an instance such that $\mathcal{Q}_1(I_1) \neq \emptyset$ and $\mathcal{Q}_2(I_1) = \emptyset$. Thanks to Lemma 5.5, we can assume that I_1 is over a domain of size at most $\mathit{varmax}(\mathcal{Q}_1)$. Thanks to genericity, we can assume that the domain is a subset of U. Let V_1 be a valuation that satisfies query $\mathcal{Q}_1$ on I_1. Furthermore, let V_2 be some consistent valuation for $\mathcal{Q}_2$ and $I_2 \stackrel{\text{def}}{=} V_2(\mathit{pos}_{\mathcal{Q}_2})$, which exists thanks to our assumption that $\mathcal{Q}_2$ is satisfiable.

To show that $\mathcal{Q}$ is not parallel-complete under $\boldsymbol{P}$, we define $I \stackrel{\text{def}}{=} \alpha_1(I_1) \cup \alpha_2(I_2) \cup \{\texttt{Type}(1), \texttt{Start}_1(1), \texttt{Start}_2(1), \texttt{Start}_2(2)\}$. Then, valuation V which maps all variables in $\mathcal{Q}_1$ and $\mathcal{Q}_2$ as V_1 and V_2, respectively, and $\ell_1 \mapsto 1$, $\ell_2 \mapsto 2$, and $t \mapsto 1$ satisfies query $\mathcal{Q}$ on instance I. However, there is no locally satisfying valuation for $\mathcal{Q}$. For a contradiction, assume existence of such a valuation W. Since $\texttt{Type}(1)$ and $\texttt{Start}_1(1)$ are the only $\texttt{Type}$- and $\texttt{Start}_1$-facts contained in the global instance, we have $W(t) = W(\ell_1) = 1$. By definition of $\boldsymbol{P}$, this valuation can only be satisfying on node κ_1. Since only $\texttt{Start}_2(1) \in \mathit{loc\text{-}inst}_{\boldsymbol{P},I}(\kappa_1)$, Claim 10.1.1 implies existence of a satisfying valuation W_2 for $\mathcal{Q}_2$ on $\alpha_1^{-1}(I) = I_1$, which contradicts the choice of instance I_1. Hence, query $\mathcal{Q}$ is not parallel-complete under policy $\boldsymbol{P}$.

To show that $\mathcal{Q}$ is also not parallel-sound under $\boldsymbol{P}$, let instance $I \stackrel{\text{def}}{=} \alpha_1(I_1) \cup \alpha_2(I_2) \cup \{\texttt{Type}(2), \texttt{Start}_1(1), \texttt{Start}_2(1), \texttt{Start}_2(2), \texttt{Stop}(2)\}$. Then, valuation V which maps all variables in $\mathcal{Q}_1$ and $\mathcal{Q}_2$ as V_1 and V_2, respectively, and $\ell_1 \mapsto 1$, $\ell_2 \mapsto 2$, and $t \mapsto 2$ satisfies query $\mathcal{Q}$ on the local instance $\mathit{loc\text{-}inst}_{\boldsymbol{P},I}(\sigma_1)$ of node σ_1 because this node is not responsible for fact $\texttt{Stop}(2)$. However, there is no globally satisfying valuation for $\mathcal{Q}$. Towards a contradiction, assume existence of such a valuation W. Since $\texttt{Start}_2(1), \texttt{Start}_2(2)$ are the only $\texttt{Start}_2$-facts in the global instance, we have $W(\ell_2) = 1$ or $W(\ell_2) = 2$. The latter cannot hold because the valuation then prohibits the present fact $\texttt{Stop}(2)$. The former implies, by Claim 10.1.1, a satisfying valuation W_2 for $\mathcal{Q}_2$ on $\alpha_1^{-1}(I) = I_1$, which contradicts the choice of instance I_1. Hence, query $\mathcal{Q}$ is not parallel-sound under policy $\boldsymbol{P}$.

This completes the proof that problem Containment($\mathbf{BCQ}^\neg, \mathbf{BCQ}^\neg$) is reducible to all three problems Parallel-Complete($\mathbf{CQ}^\neg, \mathcal{P}_{\text{rule}}$), Parallel-Sound($\mathbf{CQ}^\neg, \mathcal{P}_{\text{rule}}$), and Parallel-Correct($\mathbf{CQ}^\neg, \mathcal{P}_{\text{rule}}$) in polynomial time and thus shows coNEXPTIME-hardness via Corollary 5.3. □

11 Proofs for Section 7: Full conjunctive queries

In the proofs below we drop the convention that in UCQs all disjuncts have pairwise disjoint variable sets (which was introduced to simplify the proofs in

other sections). However, one can easily observe that a UCQ that does not comply with this convention can be easily transformed to one that does, by for example, adding distinct indices to variables in separate disjuncts.

11.1 Proof for Theorem 7.1

Theorem 7.1 follows from the hardness results in Proposition 11.1, and the upper bound in Proposition 11.3, which are given below.

Proposition 11.1. 1. CONTAINMENT($\mathbf{FCQ}^{\neg,\neq}, \mathbf{FCQ}^{\neg,\neq}$) is in P;

2. CONTAINMENT($\mathbf{FCQ}^{\neg,\neq}, \mathbf{UFCQ}^{\neg,\neq}$) is coNP-hard.

Proof. **1.** Let $\mathcal{Q}_1, \mathcal{Q}_2 \in \mathbf{FCQ}^{\neg,\neq}$. We show that $\mathcal{Q}_1 \subseteq \mathcal{Q}_2$ if and only if $h(pos_{\mathcal{Q}_2}) \subseteq pos_{\mathcal{Q}_1}$ and $h(neg_{\mathcal{Q}_2}) \subseteq neg_{\mathcal{Q}_1}$ for the substitution h that identifies the head relation,[9] $h(head_{\mathcal{Q}_2}) = head_{\mathcal{Q}_1}$. It then immediately follows that CONTAINMENT($\mathbf{FCQ}^{\neg,\neq}, \mathbf{FCQ}^{\neg,\neq}$) is in P;

To show the claim, let $\mathcal{Q}_1 \subseteq \mathcal{Q}_2$. Let I^- be the minimal canonical database for $\mathcal{Q}_1$, i.e., I^- consists of the frozen atoms in $pos_{\mathcal{Q}_1}$. By I^+ we denote the maximal canonical database for $\mathcal{Q}_1$, i.e., I^+ contains every frozen atom over $vars(\mathcal{Q})$ (and the relations where $\mathcal{Q}_1$ and $\mathcal{Q}_2$ are defined over) that is not in $neg_{\mathcal{Q}_1}$. By construction, $head_{\mathcal{Q}_1} \in \mathcal{Q}_1(I^-)$ and $head_{\mathcal{Q}_1} \in \mathcal{Q}_1(I^+)$, which implies by containment that $head_{\mathcal{Q}_1} \in \mathcal{Q}_2(I^-)$ and $head_{\mathcal{Q}_1} \in \mathcal{Q}_2(I^+)$. By fullness of $\mathcal{Q}_2$, both are derived by the same valuation V. Now, the former implies $V(pos_{\mathcal{Q}_2}) \subseteq I^- = pos_{\mathcal{Q}_1}$, and the latter implies $V(neg_{\mathcal{Q}_2}) \cap I^+ = \emptyset$. Thus, $V(neg_{\mathcal{Q}_2}) \subseteq neg_{\mathcal{Q}_1}$. Hence, V describes the desired substitution.

For the other direction, suppose that substitution h has the desired properties. Let I be an arbitrary instance, and $\boldsymbol{f} \in \mathcal{Q}_1(I)$. Thus, there is a valuation V_1, where $V_1(head_{\mathcal{Q}_1}) = \boldsymbol{f}$, $V_1(pos_{\mathcal{Q}_1}) \subseteq I$, and $V_1(neg_{\mathcal{Q}_1}) \cap I = \emptyset$. Let $V_2 \stackrel{\text{def}}{=} V_1 \circ h$. Then, $V_2(head_{\mathcal{Q}_2}) = V_1(head_{\mathcal{Q}_1}) = \boldsymbol{f}$, $V_2(pos_{\mathcal{Q}_2}) \subseteq V_1(pos_{\mathcal{Q}_1}) \subseteq I$, and $V_2(neg_{\mathcal{Q}_2}) \subseteq V_1(neg_{\mathcal{Q}_1})$. The latter implies $V_2(neg_{\mathcal{Q}_2}) \cap I = \emptyset$, and thus $\boldsymbol{f} \in \mathcal{Q}_2(I)$.

2. The reduction is from the graph 3-colorability problem, which is well-known to be NP-complete, and asks for a given graph G whether there is a coloring of the nodes in G, using only 3 colors, such that adjacent nodes have distinct colors. Let G be an arbitrary input for the described problem.

We construct queries $\mathcal{Q}_1 \in \mathbf{FCQ}^{\neg,\neq}$ and $\mathcal{Q}_2 \in \mathbf{UFCQ}^{\neg,\neq}$, and show $\mathcal{Q}_1 \subseteq \mathcal{Q}_2$ if and only if G is *not* 3-colorable. Intuitively, this is done by letting $\mathcal{Q}_1$ derive tuples that each represent a complete labelling function for (a substitution of) graph G. Semantically, query $\mathcal{Q}_2$ is very similar to $\mathcal{Q}_1$, but derives only invalid labelling functions, that is, which either give two adjacent nodes the same color, or use more than three colors. The latter is implemented as a union of queries where each disjunct detects a particular issue.

[9] In particular, containment implies existence of such a substitution despite possible multiple occurences of the same variable in $head_{\mathcal{Q}_2}$, as the following proof shows.

Queries $\mathcal{Q}_1$ and $\mathcal{Q}_2$ are defined over database schema $\mathcal{D} \stackrel{\text{def}}{=} \{\mathtt{E}^{(2)}, \mathtt{L}^{(2)}\}$, where $\mathtt{E}$ represents the edges of G, and $\mathtt{L}$ is used to model mappings from nodes in G onto colors.

Before going to the construction itself, we first define EDGES as the set of atoms describing the edges of G. We do this by taking for every node in G a unique variable in **var**, and adding the atom $\mathtt{E}(x, y)$ to EDGES if and only if the nodes represented by x and y are adjacent in G. Second, we define the set LABELS, consisting of atoms $\mathtt{L}(x, y_x)$ for every node-representing variable x, where y_x denotes a fresh variable. We call y_x a color-representing variable.

Henceforth, we denote by $\boldsymbol{x}$ the node-representing variables, and by $\boldsymbol{y}$ the color-representing variables. Both are assumed to have a fixed order.

Now, we define query $\mathcal{Q}_1$ as:

$$H(\boldsymbol{x}, \boldsymbol{y}) \leftarrow \text{EDGES}, \text{LABELS}.$$

Notice that, on given instance I, $\mathcal{Q}_1$ outputs every possible labelling function for substitutions of graph G representable with the facts in I.

Query $\mathcal{Q}_2$ is slightly more complex, therefore we chop $\mathcal{Q}_2$ down in two semantically-meaningful UFCQ$^\neg$s: $\mathcal{Q}^{\forall h}$, which detects whether a described labelling function assigns the same color to some pair of adjacent nodes; and $\mathcal{Q}^{>3}$, which detects whether a described labelling function uses more than three colors.

We describe the sketched queries in more detail. To this end, let $x_1, x_2 \in \boldsymbol{x}$ be two distinct variables, representing adjacent nodes in G. Based on x_1 and x_2, we can define a substitution h mapping all the variables of $\boldsymbol{x}$ and $\boldsymbol{y}$ onto themselves, except for y_{x_1} and y_{x_2}, which are mapped onto a fresh variable y. We call h a *collision-revealing substitution* for $\boldsymbol{x}$ and $\boldsymbol{y}$.

Now $\mathcal{Q}^{\forall h}$ is the query defined as the union of queries:

$$h(\mathit{head}_{\mathcal{Q}_1}) \leftarrow h(\mathit{body}_{\mathcal{Q}_1}),$$

for all collision-revealing substitutions h for $\boldsymbol{x}$ and $\boldsymbol{y}$ as described above.

Query $\mathcal{Q}^{\forall h}$ outputs exactly those labelling functions outputted by $\mathcal{Q}_1$ that assign a same color to at least some pair of adjacent nodes.

Query $\mathcal{Q}^{>3}$ is defined as the union of queries:

$$\mathit{head}_{\mathcal{Q}_1} \leftarrow \mathit{body}_{\mathcal{Q}_1}, \forall_{i<j} y_i \neq y_j,$$

for every combination of distinct variables y_1, y_2, y_3, y_4 in $\boldsymbol{y}$.

As the construction of the individual FCQ$^\neg$ can be done in polynomial time in the size of G, and there are only n^2 many collision-revealing substitutions, the construction of $\mathcal{Q}_1$ and $\mathcal{Q}_2$ can be done in polynomial time in the size of G.

Correctness. We show that $\mathcal{Q}_1 \not\subseteq \mathcal{Q}_2$ if and only if G is 3-colorable.

(If) Suppose G is 3-colorable. So, there is a labelling function ℓ mapping the nodes in G onto 3-colors, such that no two adjacent nodes are assigned the same color. We abuse notation and assume nodes in G, as well as the colors in

the image of ℓ, are over **dom**. Let $I \stackrel{\text{def}}{=} \{\mathtt{E}(n_1, n_2) \mid (n_1, n_2) \text{ is an edge in } G\} \cup \{\mathtt{L}\big(n, \ell(n)\big) \mid n \text{ is a node in } G\}$.

Now, let V_1 be the valuation mapping the node-representing variables in $\mathcal{Q}_1$ onto the nodes of G that they represent, and the color-representing variables onto colors as defined by ℓ. Obviously, V_1 satisfies on I and derives a fact $\boldsymbol{f}$.

By choice of ℓ, $\boldsymbol{f} \notin \mathcal{Q}^{\forall h}(I)$, because none of its adjacent nodes are labelled the same color. Further, ℓ introduces only three colors, and thus $\boldsymbol{f} \notin \mathcal{Q}^{>3}(I)$, which implies $\boldsymbol{f} \notin \mathcal{Q}_2(I)$. Hence, $\mathcal{Q}_1 \not\subseteq \mathcal{Q}_2$.

(Only if) Suppose $\mathcal{Q}_1 \not\subseteq \mathcal{Q}_2$. Thus, there is an instance I and fact $\boldsymbol{f}$, where $\boldsymbol{f} \in \mathcal{Q}_1(I)$, and $\boldsymbol{f} \notin \mathcal{Q}_2(I)$. Let V_1 be the valuation for $\mathcal{Q}_1$ on I.

The former implies that I contains an interpretation of the described graph. Notice that $\boldsymbol{f}$ does not necessarily describe G itself, but rather a substitution of G under which some nodes might have been collapsed. We consider the labelling function ℓ, defined as $\ell(n) \stackrel{\text{def}}{=} V_1(x)$, for every $x \in \boldsymbol{x}$, where x is the variable representing node n in G.

By fullness of the considered queries, for each of the disjuncts of $\mathcal{Q}_2$, there is only one valuation that can possibly derive $\boldsymbol{f}$, and $\boldsymbol{f}$ uniquely defines this valuation for the respective disjunct. Therefore, from $\boldsymbol{f} \notin \mathcal{Q}^{\forall h}(I)$ we directly obtain that for every pair of adjacent nodes n_1, n_2 in G: $\ell(n_1) \neq \ell(n_2)$. So, ℓ describes a valid labelling function for G. It remains to argue that ℓ does not use more than three colors. The latter follows from $\boldsymbol{f} \notin \mathcal{Q}^{>3}(I)$, and thus the substitution of G described by $\boldsymbol{f}$ must be 3-colorable. From this it immediately follows that G must be 3-colorable as well. □

We next show that the problem remains coNP-hard, even for queries without inequalities.

Lemma 11.2. $\textsc{Containment}(\mathbf{FCQ}^{\neg,\neq}, \mathbf{UFCQ}^{\neg,\neq}) \leq_p \textsc{Containment}(\mathbf{FCQ}^{\neg}, \mathbf{UFCQ}^{\neg})$.

Proof. For this to see, let $\mathcal{Q}_1 \in \mathbf{FCQ}^{\neg,\neq}$ and $\mathcal{Q}_2 \in \mathbf{UFCQ}^{\neg,\neq}$ over some schema $\mathcal{D}$. Let $\boldsymbol{y}$ represent the variables used in $\mathcal{Q}_1$, in some fixed order. We construct a query $\mathcal{Q}_2' \in \mathbf{UFCQ}^{\neg}$ such that $\mathcal{Q}_1 \subseteq \mathcal{Q}_2$ if and only if $\mathcal{Q}_1 \subseteq \mathcal{Q}_2'$. For this, we extend $\mathcal{D}$ to a schema $\mathcal{D}'$ that also defines relations $\mathtt{Eq}^{(2)}$ and $\mathtt{Neq}^{(2)}$. Intuitively, $\mathtt{Eq}^{(2)}$ models the equality relation $=$, and $\mathtt{Neq}$ models the inequality relation $\neq$. We assume $\mathtt{Eq}$ and $\mathtt{Neq}$ are not in $\mathcal{D}$.

For the construction, we divide $\mathcal{Q}_2'$ into two subqueries $\mathcal{Q}_2^*$ and $\mathcal{Q}^{\mathtt{Neq}}$. Query $\mathcal{Q}_2^*$ results from $\mathcal{Q}_2$ by replacing every inequality $x \neq y$ by atom $\mathtt{Eq}(x, y)$. The purpose of query $\mathcal{Q}^{\mathtt{Neq}}$ is to allow derivation of a fact derivable by $\mathcal{Q}_1$ on every instance where $\mathtt{Eq}$ and $\mathtt{Neq}$ do not represent equality relation $=$ or inequality relation $\neq$, respectively.

To this end, we further divide $\mathcal{Q}^{\mathtt{Neq}}$ into the following queries: $\mathcal{Q}^{\mathtt{Neq}}_{\text{equal}}$, which detects values that are wrongly identified as being unequal (by $\mathtt{Neq}$); $\mathcal{Q}^{\mathtt{Eq},\mathtt{Neq}}_{\text{amb}}$, which detects that some values occur in both $\mathtt{Eq}$ and $\mathtt{Neq}$; $\mathcal{Q}^{\mathtt{Eq}}_{\text{asym}}$, which detects values where $\mathtt{Neq}$ is not symmetric for; $\mathcal{Q}^{\mathtt{Neq}}_{\text{asym}}$, which detects values where $\mathtt{Eq}$ is not symmetric for; $\mathcal{Q}^{\mathtt{Eq},\mathtt{Neq}}_{\text{undef}}$, which detects that certain values are not in $\mathtt{Neq}$, nor

in $\texttt{Eq}$; and finally, $\mathcal{Q}^{\texttt{Neq}}$, in which all occurrences of $x \neq y$ are replaced by atoms $\texttt{Neq}(x, y)$.

More formally, $\mathcal{Q}^{\texttt{Neq}}_{\text{equal}}$ is defined as the union of the queries:

$$h(\mathit{head}_{\mathcal{Q}_1}) \leftarrow h(\mathit{body}_{\mathcal{Q}_1}), \texttt{Neq}(y, y),$$

for all collision-revealing substitutions h for $\boldsymbol{x}$ and $\boldsymbol{y}$.

Thus, $\mathcal{Q}^{\texttt{Neq}}_{\text{equal}}$ outputs on a given instance exactly those facts outputed by $\mathcal{Q}_1$ in which some values are wrongly identified as being inequal.

We define query $\mathcal{Q}^{\texttt{Eq},\texttt{Neq}}_{\text{amb}}$, as the union over queries:

$$\mathit{head}_{\mathcal{Q}_1} \leftarrow \mathit{body}_{\mathcal{Q}_1}, \texttt{Eq}(y_1, y_2), \texttt{Neq}(y_1, y_2),$$

for every two distinct variables $y_1, y_2 \in \boldsymbol{y}$. Query $\mathcal{Q}^{\texttt{Eq}}_{\text{asym}}$ is defined as the union over queries:

$$\mathit{head}_{\mathcal{Q}_1} \leftarrow \mathit{body}_{\mathcal{Q}_1}, \texttt{Eq}(y_1, y_2), \neg\texttt{Eq}(y_2, y_1),$$

for every two distinct variables $y_1, y_2 \in \boldsymbol{y}$. Query $\mathcal{Q}^{\texttt{Neq}}_{\text{asym}}$ is defined as the union over queries:

$$\mathit{head}_{\mathcal{Q}_1} \leftarrow \mathit{body}_{\mathcal{Q}_1}, \texttt{Neq}(y_1, y_2), \neg\texttt{Neq}(y_2, y_1).$$

for every two distinct variables $y_1, y_2 \in \boldsymbol{y}$. And, eventually, query $\mathcal{Q}^{\texttt{Eq},\texttt{Neq}}_{\text{undef}}$ is defined as the union over queries:

$$\mathit{head}_{\mathcal{Q}_1} \leftarrow \mathit{body}_{\mathcal{Q}_1}, \neg\texttt{Eq}(y_1, y_2), \neg\texttt{Neq}(y_1, y_2),$$

for every two distinct variables $y_1, y_2 \in \boldsymbol{y}$.

Correctness. It remains to show $\mathcal{Q}_1 \subseteq \mathcal{Q}_2$ if and only if $\mathcal{Q}_1 \subseteq \mathcal{Q}'_2$.

(If) Suppose $\mathcal{Q}_1 \not\subseteq \mathcal{Q}_2$. Thus, for some instance I and fact $\boldsymbol{f}$, $\boldsymbol{f} \in \mathcal{Q}_1(I)$, while $\boldsymbol{f} \notin \mathcal{Q}_2(I)$. We consider instance I' over $\mathcal{D}'$, which consists of all facts in I, and for every value $a \in \mathit{adom}(I)$ a fact $\texttt{Eq}(a, a)$, and for every pair of distinct values $a, b \in \mathit{adom}(I)$ a fact $\texttt{Neq}(a, b)$.

Because $\mathcal{Q}_1$ does not reference relations $\texttt{Eq}$ or $\texttt{Neq}$, it follows that $\boldsymbol{f} \in \mathcal{Q}_1(I')$. As relations $\texttt{Eq}$ and $\texttt{Neq}$ express exactly $=$ and $\neq$, over the values in $\mathit{adom}(I)$, we obtain $\mathcal{Q}_2(I) = \mathcal{Q}^*_2(I')$. Consequently, $\boldsymbol{f} \notin \mathcal{Q}_2(I)$ implies $\boldsymbol{f} \notin \mathcal{Q}^*_2(I')$. Further, as $\texttt{Eq}$ and $\texttt{Neq}$ are symmetric, consistent, and completely defined over $\mathit{adom}(I')$, $\mathcal{Q}^{\texttt{Neq}}(I') = \emptyset$. Hence, $\boldsymbol{f} \notin \mathcal{Q}'_2(I')$.

(Only if) Suppose $\mathcal{Q}_1 \not\subseteq \mathcal{Q}'_2$. Thus, for some instance I and fact $\boldsymbol{f}$, $\boldsymbol{f} \in \mathcal{Q}_1(I)$, and $\boldsymbol{f} \notin \mathcal{Q}'_2(I)$. Let V be the valuation deriving $\boldsymbol{f}$ for $\mathcal{Q}_1$ on I and $D \stackrel{\text{def}}{=} \mathit{adom}(V(\mathit{body}_{\mathcal{Q}_1}))$. From $\boldsymbol{f} \notin \mathcal{Q}'_2(I)$ it now follows that $\texttt{Eq}$ and $\texttt{Neq}$ are well-defined over D, that is, the relations $\texttt{Eq}$ and $\texttt{Neq}$ are identical to $=$ and $\neq$ over values in D. Indeed, for all $a, b \in D$:

- there are no facts $\texttt{Neq}(a, a)$ (because $\boldsymbol{f} \notin \mathcal{Q}^{\texttt{Neq}}_{\text{equal}}(I)$);

- $\texttt{Eq}(a,b)$ in I implies $\texttt{Eq}(b,a)$ (because $\boldsymbol{f} \notin \mathcal{Q}^{\texttt{Neq}}_{\text{asym}}(I)$);
- $\texttt{Neq}(a,b)$ in I implies $\texttt{Neq}(b,a)$ (because $\boldsymbol{f} \notin \mathcal{Q}^{\texttt{Eq}}_{\text{asym}}(I)$); and
- there is a fact $\texttt{Eq}(a,b)$ or $\texttt{Neq}(a,b)$ in I (from $\boldsymbol{f} \notin \mathcal{Q}^{\texttt{Eq},\texttt{Neq}}_{\text{undef}}(I)$).

As a result, we can simply replace every occurrence of $\texttt{Neq}(x,y)$ in query $\mathcal{Q}_2^*$ by $x \neq y$ without changing its semantics over D. More formally, we obtain $\mathcal{Q}_2^*(I_{|D}) = \mathcal{Q}_2(I_{|D})$.

By fullness of $\mathcal{Q}_2$ and the definition of D, $\boldsymbol{f} \in \mathcal{Q}_2(I)$ would imply $\boldsymbol{f} \in \mathcal{Q}_2(I_{|D})$ and thus $\boldsymbol{f} \in \mathcal{Q}_2'(I)$. Consequently, it must be that $\boldsymbol{f} \notin \mathcal{Q}_2(I)$. □

Proposition 11.3. Containment($\mathbf{UFCQ}^{\neg,\neq}, \mathbf{UFCQ}^{\neg,\neq}$) is in coNP.

Proof. We observe that when $\mathcal{Q} \not\subseteq \mathcal{Q}'$, there is an instance I and fact $\boldsymbol{f}$, such that $\boldsymbol{f} \in \mathcal{Q}(I)$, and $\boldsymbol{f} \notin \mathcal{Q}'(I)$. In particular, by fullness of $\mathcal{Q}_1$ and $\mathcal{Q}_2$, there is such an instance of size at most $\ell \stackrel{\text{def}}{=} \max_{i\in\{1,\ldots,n\}}\{|pos_{\mathcal{Q}_i}|\} + m$, where n denotes the number of disjuncts of $\mathcal{Q}$, and m the number of disjuncts in $\mathcal{Q}'$.

Indeed, to see this, let V and $\mathcal{Q}_i$ be the valuation and disjunct of $\mathcal{Q}$ where $\boldsymbol{f}$ is derived by on I. Let $J \stackrel{\text{def}}{=} V(pos_{\mathcal{Q}_i})$. By fullness it follows that for each disjunct $\mathcal{Q}_j$ of $\mathcal{Q}'$ there is at most one valuation V_j eligible to derive $\boldsymbol{f}$ for $\mathcal{Q}_j$. As, by choice of $\boldsymbol{f}$, V_j does not satisfy on I it must be that either $V_j(pos_{\mathcal{Q}_j}) \not\subseteq I$, or $V_j(neg_{\mathcal{Q}_j}) \cap I \neq \emptyset$. We ignore the former. In the latter case we choose one fact from $V_j(neg_{\mathcal{Q}_j}) \cap I$ and add it to J. One can now easily verify that J is as desired.

As the above shows that there always is a witnessing instance I of size ℓ, given I, V, and $\mathcal{Q}_i$ as polynomial size certificate, we can easily verify that $\mathcal{Q}$ is indeed not contained in $\mathcal{Q}'$, by simply verifying that V indeed satisfies for disjunct $\mathcal{Q}_i$ of $\mathcal{Q}$ on I, and for all m eligible valuations for conjuncts of $\mathcal{Q}'$, either not all required facts are in I, or at least one of the prohibited facts is present. □

11.2 Proof of Proposition 7.2

The result follows from Proposition 11.1 and the reductions below.

Proposition 11.4. The following reducibility relations hold, for every $\mathcal{P} \in \{\mathcal{P}_{\text{rule}}\} \cup \mathfrak{P}_{\text{npoly}}$:

1. Containment($\mathbf{FCQ}^{\neg}, \mathbf{UFCQ}^{\neg}$) $\leq_p$ Parallel-Sound($\mathbf{UFCQ}^{\neg}, \mathcal{P}$);
2. Containment($\mathbf{FCQ}^{\neg}, \mathbf{UFCQ}^{\neg}$) $\leq_p$ Parallel-Complete($\mathbf{UFCQ}^{\neg}, \mathcal{P}$);
3. Containment($\mathbf{FCQ}^{\neg}, \mathbf{UFCQ}^{\neg}$) $\leq_p$ Parallel-Correct($\mathbf{UFCQ}^{\neg}, \mathcal{P}$).

Proof. The idea underlying the following reductions is simple: extend both CQ$^{\neg}$s $\mathcal{Q}$ and $\mathcal{Q}'$ for the containment problem by a nullary atom $\texttt{Global}()$ or its negation $\neg\texttt{Global}()$, respectively, and combine them (by union) into a single query $\mathcal{Q}^*$. By mapping the fact $\texttt{Global}()$ onto an isolated node, the distribution

policy then allows to control global and local derivability on behalf of $\mathcal{Q}$ and $\mathcal{Q}'$. Without loss of generality, we always assume that queries $\mathcal{Q}$ and $\mathcal{Q}'$ do not use the auxiliary relation `Global`.

For a CQ$^\neg$ $\mathcal{Q}$, let $\mathcal{Q}_{\texttt{Global}}$ and $\mathcal{Q}_{\neg\texttt{Global}}$ denote the queries obtained by adding the literal `Global`() or $\neg$`Global`() to $\mathcal{Q}$, respectively. For unions of CQ$^\neg$s this particularly means adding `Global`() or $\neg$`Global`() to every disjunct of $\mathcal{Q}$. The following identities can be easily proven to hold for every query $\mathcal{Q}$ and every instance I:

$$\mathcal{Q}_{\texttt{Global}}(I \cup \{\texttt{Global}()\}) = \mathcal{Q}(I); \tag{1a}$$
$$\mathcal{Q}_{\texttt{Global}}(I \setminus \{\texttt{Global}()\}) = \emptyset; \tag{1b}$$
$$\mathcal{Q}_{\neg\texttt{Global}}(I \cup \{\texttt{Global}()\}) = \emptyset; \tag{1c}$$
$$\mathcal{Q}_{\neg\texttt{Global}}(I \setminus \{\texttt{Global}()\}) = \mathcal{Q}(I). \tag{1d}$$

We only argue the reductions for policy class $\mathcal{P}_{\text{rule}}$. It is obvious that such a policy can also be represented by a policy from any class $\mathcal{P}^k_{\text{npoly}} \in \mathfrak{P}_{\text{npoly}}$.

1. We start with Containment($\mathbf{FCQ}^\neg, \mathbf{FCQ}^\neg$) $\leq_p$ Parallel-Sound($\mathbf{UFCQ}^\neg, \mathcal{P}_{\text{rule}}$).

Let $\mathcal{Q}$ and $\mathcal{Q}'$ be CQ$^\neg$s. We define a UCQ$^\neg$ $\mathcal{Q}^*$ and a policy $\boldsymbol{P}$ as follows. For this, let $\mathcal{Q}^* \stackrel{\text{def}}{=} \mathcal{Q}_{\neg\texttt{Global}} \cup \mathcal{Q}'_{\texttt{Global}}$.

We construct a subset D of $\mathbf{dom}$, where $|D| = |\mathit{vars}(\mathcal{Q}_1)|$. Since the actual data values do not matter, we can choose those with the shortest representation length and thus also represent set D polynomially in the size of query $\mathcal{Q}_1$. Now, $\boldsymbol{P}$ is defined as a distribution policy over network $\mathcal{N} = \{\kappa_1, \kappa_2\}$ that forwards every fact over D except `Global`() to κ_1, and `Global`() to node κ_2. As the described distribution policy can be straightforwardly expressed with a distribution policy in $\mathcal{P}_{\text{fin}}$, the construction of both $\mathcal{Q}^*$ and $\boldsymbol{P}$ can be done in polynomial time.

Correctness. It remains to show that $\mathcal{Q} \subseteq \mathcal{Q}'$ if and only if $\mathcal{Q}^*$ is parallel-sound under $\boldsymbol{P}$. The following observations are crucial for the correctness argument.

First of all, for each instance I that contains `Global`(), $\mathcal{Q}^*$ is equivalent to $\mathcal{Q}'$:

$$\mathcal{Q}^*(I) = \mathcal{Q}_{\neg\texttt{Global}}(I) \cup \mathcal{Q}'_{\texttt{Global}}(I) = \emptyset \cup \mathcal{Q}'(I) = \mathcal{Q}'(I). \tag{2}$$

Second, for each instance I that does not contain `Global`(), $\mathcal{Q}^*$ is equivalent to $\mathcal{Q}$:

$$\mathcal{Q}^*(I) = \mathcal{Q}_{\neg\texttt{Global}}(I) \cup \mathcal{Q}'_{\texttt{Global}}(I) = \mathcal{Q}(I) \cup \emptyset = \mathcal{Q}(I) = \mathcal{Q}\big(I \cup \{\texttt{Global}()\}\big). \tag{3}$$

Further, as κ_2 can only contain the fact `Global`(), it follows that for each instance I we have $\mathcal{Q}(\mathit{loc\text{-}inst}_{\boldsymbol{P},I}(\kappa_2)) = \emptyset$.

(Only if) Assume $\mathcal{Q} \subseteq \mathcal{Q}'$. Let I be an arbitrary subset of $\mathit{facts}(\boldsymbol{P})$. If `Global`() $\notin I$, the local instance of node κ_1 is identical to the global instance, that is $\mathit{loc\text{-}inst}_{\boldsymbol{P},I}(\kappa_1) = I$. Now, by definition of $\boldsymbol{P}$ and thus also the result sets, $\mathcal{Q}^*\big(\mathit{loc\text{-}inst}_{\boldsymbol{P},I}(\kappa_1)\big) = \mathcal{Q}^*(I)$. In particular, this implies $\mathcal{Q}^*\big(\mathit{loc\text{-}inst}_{\boldsymbol{P},I}(\kappa_1)\big) \subseteq \mathcal{Q}^*(I)$, that is, parallel-soundness of $\mathcal{Q}^*$ under $\boldsymbol{P}$ on instance I.

If $\texttt{Global}() \in I$, the local instance is $\mathit{loc\text{-}inst}_{\boldsymbol{P},I}(\kappa_1) = I \setminus \{\texttt{Global}()\}$. By Equations (2) and (3), $\mathcal{Q}^*\big(\mathit{loc\text{-}inst}_{\boldsymbol{P},I}(\kappa_1)\big) \subseteq \mathcal{Q}^*(I)$ if and only if $\mathcal{Q}(I) \subseteq \mathcal{Q}'(I)$, which holds by assumption of containment. Therefore, in both cases query $\mathcal{Q}^*$ is parallel-sound under policy $\boldsymbol{P}$.

(If) For a proof by contraposition assume $\mathcal{Q} \not\subseteq \mathcal{Q}'$. This implies existence of an instance I where $\mathcal{Q}(I) \not\subseteq \mathcal{Q}'(I)$. Without loss of generality we can assume $\mathit{adom}(I) \subseteq D$. The latter is a safe assumption as from Lemma 5.5 it follows that an instance $J \subseteq I$ exists that preserves the desired property, and where $|\mathit{adom}(J)| \leq |D|$. Further, by genericity of $\mathcal{Q}_1$ and $\mathcal{Q}_2$, we can uniquely rename data values in J to data values in D, resulting in an instance with the desired properties.

Additionally, we may safely assume that $\texttt{Global}() \in I$ because neither $\mathcal{Q}$ nor $\mathcal{Q}'$ refers to relation `Global`. This results in a local instance $\mathit{loc\text{-}inst}_{\boldsymbol{P},I}(\kappa_1) = I \setminus \{\texttt{Global}()\}$. Again, by Equations (2) and (3), we conclude $\mathcal{Q}^*\big(\mathit{loc\text{-}inst}_{\boldsymbol{P},I}(\kappa_1)\big) = \mathcal{Q}(I) \not\subseteq \mathcal{Q}'(I) = \mathcal{Q}^*(I)$, that is $\mathcal{Q}^*$ is not parallel-sound under $\boldsymbol{P}$ on instance I. Therefore, by contraposition, parallel-soundness of query $\mathcal{Q}^*$ under policy $\boldsymbol{P}$ and domain D implies containment $\mathcal{Q} \subseteq \mathcal{Q}'$.

2. For $\textsc{Containment}(\mathbf{FCQ}^\neg, \mathbf{FCQ}^\neg) \leq_p \textsc{Parallel-Complete}(\mathbf{UFCQ}^\neg, \mathcal{P}_{\text{rule}})$ the proof is analogous. Policy $\boldsymbol{P}$ is defined as before, while query $\mathcal{Q}^* \stackrel{\text{def}}{=} \mathcal{Q}_{\texttt{Global}} \cup \mathcal{Q}'_{\neg\texttt{Global}}$, i.e. atoms $\texttt{Global}()$ and $\neg\texttt{Global}()$ are swapped between the CQ$^\neg$s compared to the soundness reduction.

Analogously to Equations (2) and (3), this leads to $\mathcal{Q}^*(I) = \mathcal{Q}(I)$ and $\mathcal{Q}^*(I) = \mathcal{Q}'(I \cup \{\texttt{Global}()\})$, for every instance I where $\texttt{Global}() \in I$ and $\texttt{Global}() \notin I$, respectively. Correctness then follows again from Lemma 5.5.

3. For $\textsc{Containment}(\mathbf{FCQ}^\neg, \mathbf{FCQ}^\neg) \leq_p \textsc{Parallel-Correct}(\mathbf{UFCQ}^\neg, \mathcal{P}_{\text{rule}})$ is also similar. Policy $\boldsymbol{P}$ is defined as before while query $\mathcal{Q}^* \stackrel{\text{def}}{=} \mathcal{Q}_{\neg\texttt{Global}} \cup \mathcal{Q}'$. If $\texttt{Global}() \in I$ then $\mathcal{Q}^*(I) = \mathcal{Q}'(I)$ just as in the first reduction. However, if $\texttt{Global}() \notin I$ then $\mathcal{Q}^*(I) = \mathcal{Q}(I) \cup \mathcal{Q}'(I)$.

(Only if) Assume $\mathcal{Q} \subseteq \mathcal{Q}'$. Thus, $\mathcal{Q} \cup \mathcal{Q}' \equiv \mathcal{Q}'$ imply $\mathcal{Q}^* \equiv \mathcal{Q}'$. Let I be an arbitrary instance such that $\mathit{adom}(I) \subseteq D$. If $\texttt{Global}() \notin I$, the local instance of node κ_1 is identical to the global instance (i.e., $\mathit{loc\text{-}inst}_{\boldsymbol{P},I}(\kappa_1) = I$) by definition of $\boldsymbol{P}$, and thus also $\mathcal{Q}^*\big(\mathit{loc\text{-}inst}_{\boldsymbol{P},I}(\kappa_1)\big) = \mathcal{Q}^*(I)$.

If $\texttt{Global}() \in I$ is $\mathit{loc\text{-}inst}_{\boldsymbol{P},I}(\kappa_1) = I \setminus \{\texttt{Global}()\}$. Thus, $\mathcal{Q}^*\big(\mathit{loc\text{-}inst}_{\boldsymbol{P},I}(\kappa_1)\big) = \mathcal{Q}'(\mathit{loc\text{-}inst}_{\boldsymbol{P},I}(\kappa_1)) = \mathcal{Q}'(I) = \mathcal{Q}^*(I)$. Therefore, in both cases query $\mathcal{Q}^*$ is parallel-correct under policy $\boldsymbol{P}$.

(If) We assume $\mathcal{Q} \not\subseteq \mathcal{Q}'$. This implies again by Lemma 5.5 and genericity of $\mathcal{Q}$ and $\mathcal{Q}'$, the existence of an instance I where $\mathcal{Q}(I) \not\subseteq \mathcal{Q}'(I)$ and $\mathit{adom}(I) \subseteq D$. It follows $\mathcal{Q}(I) \cup \mathcal{Q}'(I) \not\equiv \mathcal{Q}'(I)$. We may safely assume that $\texttt{Global}() \in I$ because neither $\mathcal{Q}$ nor $\mathcal{Q}'$ refers to relation `Global`. Therefore, $\mathit{loc\text{-}inst}_{\boldsymbol{P},I}(\kappa_1) = I \setminus \{\texttt{Global}()\}$. Again, by Equations (2) and (3), we conclude $\mathcal{Q}^*\big(\mathit{loc\text{-}inst}_{\boldsymbol{P},I}(\kappa_1)\big) = \mathcal{Q}(I) \cup \mathcal{Q}'(I) \neq \mathcal{Q}'(I) = \mathcal{Q}^*(I)$, that is, $\mathcal{Q}^*$ is not parallel-correct under $\boldsymbol{P}$ on instance I, the desired contradiction. □

11.3 Proof for Theorem 7.3

Proof. Let $\mathcal{Q}$ be a UFCQ$^{\neg,\neq}$ with m disjuncts, and $\boldsymbol{P} = (U, \mathit{rfacts}_{\boldsymbol{P}}) \in \mathcal{P}_{\text{rule}}$.

1. From the definition of parallel-soundness we immediately obtain that $\mathcal{Q}$ is not parallel-sound on $\boldsymbol{P}$ if and only if there is an instance $I \subseteq \mathit{facts}(\mathcal{D}, U)$, a fact $\boldsymbol{f}$, and a node κ, where $\boldsymbol{f} \in \mathcal{Q}(\mathit{loc\text{-}inst}_{\boldsymbol{P},I}(\kappa))$ and $\boldsymbol{f} \notin \mathcal{Q}(I)$. The former implies that there is a valuation V for some of the disjuncts of $\mathcal{Q}$, say $\mathcal{Q}_i$, which derives $\boldsymbol{f}$ on $\mathit{loc\text{-}inst}_{\boldsymbol{P},I}(\kappa)$. The latter means that for all disjuncts $\mathcal{Q}_j$, and valuations V_j for $\mathcal{Q}_j$, where $V_j(\mathit{head}_{\mathcal{Q}_j}) = \boldsymbol{f}$, V_j fails to satisfy for $\mathcal{Q}_j$ on I.

As $\mathcal{Q}$ is full, for each disjunct $\mathcal{Q}_j$ there is at most one eligible valuation V_j, which is uniquely defined by the head of $\mathcal{Q}_j$ and the desired output fact $\boldsymbol{f}$. We thus have to check only m valuations V_j. From this observation it also follows that we can restrict the above condition to instances I of size at most $\max_{j\in\{1,\ldots,m\}}\{|\mathit{pos}_{\mathcal{Q}_j}|\} + m$. Indeed, given I, κ, and valuation V (being the valuation showing $\boldsymbol{f} \in \mathcal{Q}(\mathit{loc\text{-}inst}_{\boldsymbol{P},I}(\kappa))$), obviously $V(\mathit{pos}_{\mathcal{Q}_i}) \subseteq I$, and, for every $j \in \{1,\ldots,m\}$, either $V_j(\mathit{pos}_{\mathcal{Q}_j}) \not\subseteq I$ or there are facts in $V_j(\mathit{neg}_{\mathcal{Q}_j}) \cap I$. In the latter case, let $\boldsymbol{g}_j$ be one of these facts. Now, we can construct an instance I' of the desired size by simply taking the facts in $V(\mathit{pos}_{\mathcal{Q}_i})$ and the chosen facts $\boldsymbol{g}_j$. It is easy to see that $\boldsymbol{f} \in \mathcal{Q}\big(\mathit{loc\text{-}inst}_{\boldsymbol{P},I'}(\kappa)\big)$, while $\boldsymbol{f} \notin \mathcal{Q}(I')$.

As $V(\mathit{pos}_{\mathcal{Q}_i}) \subseteq \mathit{rfacts}_{\boldsymbol{P}}(\kappa)$ and $V(\mathit{neg}_{\mathcal{Q}_i}) \cap \mathit{rfacts}_{\boldsymbol{P}}(\kappa) \cap I = \emptyset$ can be verified in polynomial time in $\boldsymbol{P}$, it follows that verifying violation of parallel-soundness can be done in polynomial time in the size of the input, by additionally providing an appropriate node κ, instance I, where $|I| \leq \max_{j\in\{1,\ldots,m\}}\{|\mathit{pos}_{\mathcal{Q}_j}|\} + m$, and integer i denoting disjunct $\mathcal{Q}_i$.

2. From the definition of parallel-completeness, it follows that $\mathcal{Q}$ is not parallel-complete on $\boldsymbol{P}$ if and only if there is an instance $I \subseteq \mathit{facts}(\mathcal{D}, U)$ and fact $\boldsymbol{f}$, such that $\boldsymbol{f} \in \mathcal{Q}(I)$, and for all nodes κ, $\boldsymbol{f} \notin \mathcal{Q}(\mathit{loc\text{-}inst}_{\boldsymbol{P},I}(\kappa))$. The former particularly implies that for some valuation V and integer $i \in \{1,\ldots,m\}$, $V(\mathit{head}_{\mathcal{Q}_i}) = \boldsymbol{f}$, $V(\mathit{pos}_{\mathcal{Q}_i}) \subseteq I$, and $V(\mathit{neg}_{\mathcal{Q}_i}) \cap I = \emptyset$. The latter implies that for all integers $j \in \{1,\ldots,m\}$, the valuation V_j defined by the head of $\mathcal{Q}_j$ and fact $\boldsymbol{f}$, does not satisfy on any of the nodes κ, that is, $V_j(\mathit{pos}_{\mathcal{Q}_j}) \not\subseteq \mathit{loc\text{-}inst}_{\boldsymbol{P},I}(\kappa)$ or $V_j(\mathit{neg}_{\mathcal{Q}_j}) \cap \mathit{loc\text{-}inst}_{\boldsymbol{P},I}(\kappa) \neq \emptyset$.

Again the size of I can be bounded, in particular, it suffices to consider only instances of size at most $\max_{j\in\{1,\ldots,m\}}\{|\mathit{pos}_{\mathcal{Q}_j}|\} + |\mathcal{N}| \cdot m$, where $\mathcal{N}$ denotes the network where $\boldsymbol{P}$ is defined over. Indeed, for every instance I as in the condition, we can construct an instance I' containing all the facts in $V(\mathit{pos}_{\mathcal{Q}_i})$ and for every node κ, and every $j \in \{1,\ldots m\}$ either add nothing (when $V_j(\mathit{pos}_{\mathcal{Q}_j}) \not\subseteq I$), or add some fact $\boldsymbol{g}_j \in V_j(\mathit{neg}_{\mathcal{Q}_j}) \cap I \cap \mathit{rfacts}_{\boldsymbol{P}}(\kappa)$.

By definition, the considered distribution policies have only polynomially many nodes in n, and verifying whether a fact is available on some node can be done in polynomial as well. Hence, the result follows.

3. As parallel-correctness for $\mathcal{Q}$ under $\boldsymbol{P}$ means $\mathcal{Q}$ being parallel-sound and parallel-complete under $\boldsymbol{P}$, it follows immediately from (1) and (2) that deciding parallel-correctness for **UFCQ**$^{\neg,\neq}$ and $\mathcal{P}_{\text{rule}}$ is in coNP. □